Proceedings of Machine Learning Research vol 144:1–12, 2021

Data-driven design of switching reference governors for brake-by-wire applications

Andrea Sassella ANDREA.SASSELLA@POLIMI.IT
Dipartimento di Elettronica e Informazione, Politecnico di Milano, Milano, Italy

Valentina Breschi VALENTINA.BRESCHI@POLIMI.IT
Dipartimento di Elettronica e Informazione, Politecnico di Milano, Milano, Italy

Simone Formentin SIMONE.FORMENTIN@POLIMI.IT
Dipartimento di Elettronica e Informazione, Politecnico di Milano, Milano, Italy

Abstract

Nowadays, data are ubiquitous in control design and data-driven approaches are in constant evolution. By following such trend, in this paper we propose an approach for the direct data-driven design of switching reference governors for nonlinear plants within a brake-by-wire application. The braking system is assumed to be pre-stabilized via a simple unknown controller attaining unsatisfactory performance in terms of output tracking and actuator effort. Hence, the reference governor is used to improve the overall closed-loop behavior, resulting into smoother maneuvering. Preliminary results on a simulation setup show the effectiveness of the proposed strategy, thus motivating further investigation on the topic.

Keywords: Hierarchical control, data-driven control, braking control

1. Introduction

Anti-locking brake systems (ABS) have become a standard active safety feature in modern vehicles, see Kiencke and Nielsen (2000a), and, as a consequence, braking control is widely studied within the control community. Approaches proposed to the tackle braking control range from heuristic threshold-based rules (see, e.g., Wellstead and Pettit (1997)), to model-based techniques, see, e.g., Martinez and Canudas-de-Wit (2007); Yi et al. (2000); Drakunov et al. (1995). The latter approaches mainly rely on a model for the wheel slip dynamics, which is strongly nonlinear and uncertain, due to the lack of reliable information on the road-tire condition. Despite several methods have been proposed to estimate models for the wheel slip dynamics, they usually require time consuming and expensive road tests to be reliable, see Canudas-de-Wit et al. (2003). To cope with this limitation of model-based strategies, in recent years, braking control techniques have been proposed that use road tests data to directly design the controller, without first identifying a model for the plant, see, e.g., Radac et al. (2017); Formentin et al. (2015, 2011). The approach proposed in Radac et al. (2017) relies on Q-learning (Watkins and Dayan (1992)), while in Formentin et al. (2015), a two-degree of freedom architecture is exploited to cope with the system nonlinearity. Instead, in Formentin et al. (2011), a data-driven controller learned via the Virtual Reference Feedback Tuning (VRFT) approach Campi et al. (2002) is paired with a data-driven nonlinear compensator Fliess and Join (2009), which, in turn, relies on a simpler ultra-local model for the braking system.

In this work, we assume that the braking system is already stabilized via a simple, yet unknown, controller (see Savaresi and Tanelli (2010)), which is not specifically designed by accounting for the two operating conditions (stable and unstable) of the wheel. Since it is known that optimal performance can be achieved at the boundaries of stability, see again Savaresi and Tanelli (2010), we propose to directly use data collected in closed-loop to design a reference governor to improve performance in terms of output tracking and reduce the effort required by the actuators. Differently to data-based approach already presented in the literature to learn reference governors (see, e.g., Liu et al. (2019); Chakrabarty et al. (2020)), in this work we present a fully *offline* technique to design a switching reference governor, that allows us to explicitly account for the two possible operating conditions of the braking system. Specifically, we rely on the approximation power of piecewise affine maps Breiman (1993), thus designing a Piecewise Linear (PWL) reference governor from data given some desired closed-loop performance, that, in turn, is dictated by a *user-defined* reference model as in Campi et al. (2002); Formentin et al. (2019, 2016). Thanks to this design choice, we propose an optimization-based approach inspired by the main idea employed for direct control design in Breschi and Formentin (2020). The advantage of this approach relies upon the possibility of retrieving both the local governors and the logic dictating switches from one mode to the other *directly from data*, thus requiring little to no knowledge on the inner loop.

The paper is structured as follows. The problem is formally stated in Section 2. The proposed data-driven approach for the direct design of switching reference governors is presented in Section 3, while Section 4 reports some simulation results showing the potential of the proposed strategy. The paper is ended by some concluding remarks.

2. Problem setting

Consider the longitudinal wheel slip control problem during longitudinal braking in a road vehicle. As done in Savaresi and Tanelli (2010), we assume the wheels of the vehicle to be decoupled and the longitudinal speed of the vehicle to be practically constant as compared to the slip dynamics. Under these assumptions, a braking system $\mathcal{G}$ can be modeled as follows:

$$\lambda(t+1) = f(\lambda(t), \mu(\lambda(t)), u(t)) \tag{1}$$

where $f : [0,1] \times \mathbb{R}^2 \to [0,1]$ is an *unknown* nonlinear function of the longitudinal slip of the wheel $\lambda(t) \in [0,1]$, the longitudinal friction coefficient $\mu(\lambda(t)) \in \mathbb{R}$, and the braking torque $u(t) \in \mathbb{R}$ [Nm]. As shown in Burckhardt (1993); Kiencke and Nielsen (2000b), the friction coefficient non-linearly depends on λ, typically as shown in Figure 1. Moreover, for constant braking torques, the system exhibits at most two equilibria, one before and one after the peak of the friction coefficient $\mu(\lambda)$ of Figure 1. Such equilibria establish two different operating regions, with the system being either locally *asymptotically stable* (before the peak) or *unstable* (after the peak).

Assume that the system is stabilized by an available low-performing (possibly unknown) controller $\mathcal{C}$, so that we can safely carry out closed-loop experiments with tracking reference $\rho(t) \in [0,1]$ and collect data in both its operating conditions.

Let the *measured* slip be corrupted by an additive zero-mean white noise $v(k)$, i.e.,

$$y(t) = \lambda(t) + v(t). \tag{2}$$

The purpose of this work is to use the available data $\mathcal{D}_T = \{\rho(t), y(t)\}_{t=1}^{T}$ to design a data-driven reference governor $\mathcal{R}$ that boosts the performance of the braking system up to a target closed-loop

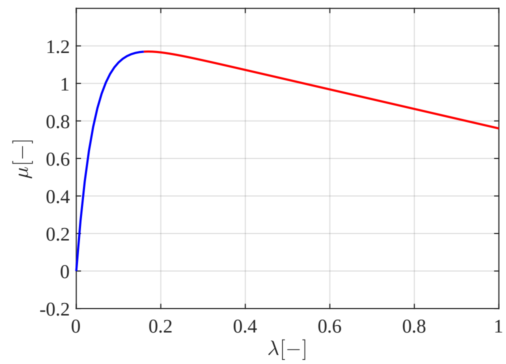

Figure 1: Longitudinal friction coefficient μ as function of the longitudinal slip λ. The blue part of the curve corresponds to an asymptotically stable linearized dynamics of (1), whereas the red part denotes the values of λ corresponding to an unstable linearized behaviour.

behavior, by differently addressing the two illustrated operating conditions. In particular, the desired closed-loop behavior is here dictated by a *user-defined* Linear Time Invariant (LTI) reference model:

$$\mathcal{M}: \begin{cases} x_M(t+1) = A_M x_M(t) + B_M r(t), \\ y_{des}(t) = C_M x_M(t) + D_M r(t), \end{cases} \tag{3}$$

where $r(t) \in [0,1]$ is the slip set-point and $y_{des}(t)$ is the corresponding output we aim at attaining in closed-loop. We stress that the reference model in (3) is chosen before the design of the reference governor $\mathcal{R}$ and, thus, it is known throughout the leaning phase.

Thus, our final braking control scheme has the hierarchical structure in Figure 2, with the data-driven reference governor $\mathcal{R}$ imposing the reference to the unknown inner-loop.

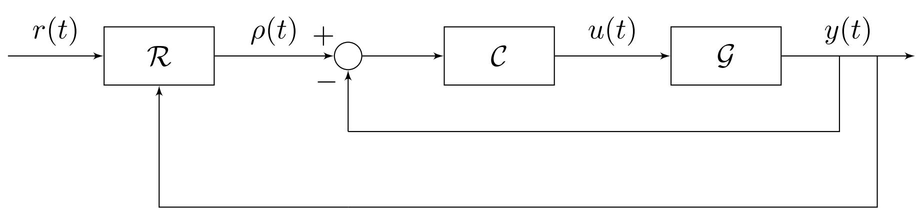

Figure 2: Hierarchical architecture of the considered control scheme.

3. Data-driven design of switching reference governors

Given the different linearized dynamics of the system (1) in the two operating conditions, according to the scheme in Figure 2 we design a *fixed-order* switching reference governor described by the following input-output model:

$$A(\Theta_{s(t)}, q)\rho(t) = B(\Theta_{s(t)}, q)(r(t) - y(t)), \tag{4a}$$

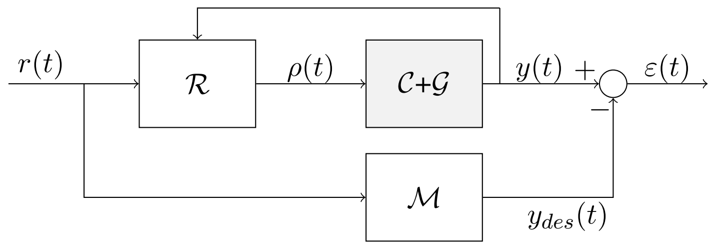

Figure 3: Data-driven method for reference governor design. The block in gray denotes the unknown inner feedback loop, including the plant $\mathcal{G}$ and the (possibly unknown) controller $\mathcal{C}$.

with $A(\Theta_{s(t)}, q)$ and $B(\Theta_{s(t)}, q)$ being the following polynomials in the shift operator q[1]:

$$A(\Theta_{s(t)}, q) = 1 + \sum_{i=1}^{n_a} \theta_i^{s(t)} q^{-i}, \qquad B(\Theta_{s(t)}, q) = \sum_{j=n_a+1}^{n_a+n_b+1} \theta_j^{s(t)} q^{-j}.$$

whose orders $n_a, n_b \in \mathbb{N}$ are fixed by the designer *a-priori*. The switching signal $s(t) \in \{1, 2, \ldots, K\}$ dictates which of the $K \in \mathbb{N}$ modes is active at time t, with K being an additional design parameter also fixed beforehand. Let $\chi(t) \in \mathcal{X} \subseteq \mathbb{R}^{n_\chi}$ be a collection of the past inputs and outputs:

$$\chi(t) = \begin{bmatrix} \rho(t-1) & \cdots & \rho(t-n_a) & y(t-1) & \cdots & y(t-n_b) \end{bmatrix}'.$$

In this work, the switching signal $s(t)$ is determined by a polyhedral partition $\{\mathcal{X}_k\}_{k=1}^K$ of the space $\mathcal{X} \subseteq \mathbb{R}^{n_\chi}$, i.e., $\mathcal{X}_k := \{\chi \text{ s.t. } \mathcal{H}_k \chi \leq \mathcal{F}_k\}$, $k = 1, \ldots, K$, and, thus, it is defined as:

$$s(t) = k \iff \chi(t) \in \mathcal{X}_k, \; k = 1, \ldots, K. \tag{4b}$$

We remark that priors on the number of possible operating conditions of $\mathcal{G}$ (like in the braking applications, where a qualitative description of the dynamics suggests $K = 2$) are likely to improve the overall performance. Nonetheless, cross-validation can be used to infer the order and the number of modes of the reference governor, whenever no priors on the inner loop are available.

The direct design of a *fixed-order* reference governor $\mathcal{R}$, like the one in (4a), thus involves: (i) learning the parameters $\Theta_k = [\, \theta_1^k \; \cdots \; \theta_{n_a+n_b+1}^k \,]'$ of each local governor (see (4a)), $k = 1, \ldots, K$, and (ii) finding a polyhedral partition $\{\mathcal{X}_k\}_{k=1}^K$ driving the switching dynamics, according to (4b).

Since the desired behavior is specified using a reference closed-loop model (3), according to the matching scheme shown in Figure 3 we want to minimize the mismatch $\varepsilon(t)$ between the actual closed-loop behavior and the desired one, namely:

$$\varepsilon(t) = y(t) - M(q)r(t), \tag{5}$$

1. $q^i u(t) = u(t+i)$ for all $i \in \mathbb{Z}$

where $M(q)$ is a shorthand for the transformation expressed by $\mathcal{M}$. Accordingly, we formulate the reference governor design problem as:

$$\begin{aligned} &\min_{\varepsilon,\{\Theta_k,\mathcal{X}_k\}_{k=1}^K} \sum_{t=1}^{T} \varepsilon^2(t) \\ &\text{s.t. } A(\Theta_{s(t)},q)\rho(t) = B(\Theta_{s(t)},q)e(t), \\ &s(t)=k \iff \chi(t)\in\mathcal{X}_k, \quad k\in\{1,\dots,K\}, \end{aligned} \tag{6}$$

where the constraints originate from the structure of the controller in (4) and they must hold for all $t \in \{1,\dots,T\}$, with $e(t) = r(t) - y(t)$ being the *tracking error*: Despite the problem in (6) explicitly depends on the available data, the reference $r(t)$ has to be fixed beforehand to solve it, which might result into a less general governor. Inspired by Campi et al. (2002), to overcome this problem we introduce the virtual reference

$$\tilde{r}(t) = M^{-1}(q)(y(t) - \varepsilon(t)), \tag{7}$$

where $M^{-1}(q)$ is the inverse of $M(q)$, that, in turn, depends on the *unknown* mismatch error $\varepsilon(t)$. The reader is referred to Breschi and Formentin (2020) for further details on the computation of the inverse map $M^{-1}(q)$.

Let $\tilde{e}(t) = \tilde{r}(t) - y(t)$ be the *fictitious tracking error*. The minimization problem in (6) thus can be recast as

$$\begin{aligned} &\min_{\varepsilon,\{\Theta_k,\mathcal{X}_k\}_{k=1}^K} \sum_{t=1}^{T} \varepsilon^2(t) \\ &\text{s.t. } A(\Theta_{s(t)},q)\rho(t) = B(\Theta_{s(t)},q)\tilde{e}(t), \\ &s(t)=k \iff \chi(t)\in\mathcal{X}_k, \quad k\in\{1,\dots,K\}, \end{aligned} \tag{8}$$

Problem (8) only depends on the available data $\{\rho(t), y(t)\} \in \mathcal{D}_T$ and on the user-defined reference model. However, due to the features of the fictitious reference in (7), the constraint on the controller dynamics makes the problem *non-convex*. Indeed, it holds that:

$$B(\Theta_{s(t)},q)\tilde{e}(t) = B(\Theta_{s(t)},q)M^{-1}(q)y(t) - B(\Theta_{s(t)},q)M^{-1}(q)\varepsilon(t), \tag{9}$$

and, thus, the controller dynamics features a biconvex term. As explained in detail in Breschi and Formentin (2020), problem (8) can be convexified by introducing the residual

$$\varepsilon_u(\Theta_{s(t)},t) = B(\Theta_{s(t)},q^{-1})M^{-1}(q)\varepsilon(t) = B(\Theta_{s(t)},q)M^{-1}(q)y(t) - A(\Theta_{s(t)},q)\rho(t), \tag{10}$$

with $t = 1,\dots,T$, which can be minimized instead of $\{\varepsilon(t)\}_{t=1}^T$.

This yields a slightly different optimization problem, which still leads to the optimal solution achieving $\mathcal{M}$ (if this is possible):

$$\begin{aligned} &\min_{\varepsilon_u,\{\Theta_k,\mathcal{X}_k\}_{k=1}^K} \sum_{t=1}^{T} \varepsilon_u^2(\Theta_{s(t)},t) + \beta\sum_{k=1}^{K}\|\Theta_k\|_2^2 \\ &\text{s.t. } s(t)=k \iff \chi(t)\in\mathcal{X}_k, \quad k\in\{1,\dots,K\}, \end{aligned} \tag{11}$$

Algorithm 1: Reference governor design
Input: Dataset $\mathcal{D}_T$; initial sequence $\mathcal{S}^0$; $\beta, \beta_p > 0$.

1. **for** $i = 1, \dots$ **do**
   1.1. $\{\Theta_k^i\}_{k=1}^K \leftarrow \text{argmin}_{\{\Theta_k\}_{k=1}^K} J_1^{IV}(\{\Theta_k\}_{k=1}^K, \mathcal{S}^{i-1})$
   1.2. $\{\omega_k^i, \gamma_k^i\}_{k=1}^K \leftarrow \text{argmin}_{\{\omega_k \gamma_k\}_{k=1}^K} J_2(\{\omega_k, \gamma_k\}_{k=1}^K, \mathcal{S}^{i-1})$
   1.3. $\mathcal{S}^i \leftarrow \arg\min_{\mathcal{S}} J_1(\{\Theta_k^i\}_{k=1}^K, \mathcal{S}) + J_2(\{\omega_k^i, \gamma_k^i\}_{k=1}^K, \mathcal{S})$
2. **until** a (customizable) stopping criterion is satisfied.

Output: Local controller parameters $\{\Theta_k^\star\}_{k=1}^K$; PWA separator $\{\omega_k^\star, \gamma_k^\star\}_{k=1}^K$.

where Tikhonov regularization is used to better condition the problem, with $\beta > 0$ being a parameter to be selected by the user.

The optimization problem in (11) is still constrained to the switching dynamics of the system, that, in turn, depends on a polyhedral partition of space $\mathcal{X}$. To remove this constraint, in this work we search for a *Piecewise Affine* (PWA) separator of such a space, that naturally enforces a polyhedral partition on $\mathcal{X}$. Specifically, we look for a function $\phi : \mathcal{X} \to \mathbb{R}$, defined as

$$\phi(\chi) = \max_{k=1,\dots,K} \phi_k(\chi), \tag{12}$$

where $\phi_k = \omega_k' \chi - \gamma_k$, for $k = 1, \dots, K$, and the parameters $\omega_k \in \mathbb{R}^{n_\chi}$ and $\gamma_k \in \mathbb{R}$ have to be estimated from data. Accordingly, each polyhedron $\mathcal{X}_k \subset \mathcal{X}$ can be defined as

$$\mathcal{X}_k = \{\chi \in \mathcal{X} : (\omega_k)'\chi - \gamma_k > (\omega_j)'\chi - \gamma_j,\ \forall j \in \{1, \dots, K\}, j \neq k\},$$

for $k = 1, \dots, K$. Thus, we cast our data-driven reference governor problem as

$$\min_{\{\Theta_k, \omega_k, \gamma_k\}_{k=1}^K, \varepsilon_u, \mathcal{S}} J_1(\{\Theta_k\}_{k=1}^K, \mathcal{S}) + J_2(\{\omega_k, \gamma_k\}_{k=1}^K, \mathcal{S}), \tag{13a}$$

where $\mathcal{S} = \{s(t)\}_{t=1}^T$ is the sequence of active local governors. Based on (11), the first term in the cost is

$$J_1(\{\Theta_k\}_{k=1}^K, \mathcal{S}) = \sum_{t=1}^T \varepsilon_u^2(\Theta_{s(t)}, t) + \beta \sum_{k=1}^K \|\Theta_k\|_2^2, \tag{13b}$$

while, by incorporating the convex loss function introduced in Breschi et al. (2016), $J_2(\{\omega_k, \gamma_k\}_{k=1}^K, \mathcal{S})$ is given by

$$\sum_{t=1}^T \sum_{\substack{j=1 \\ j \neq s(t)}}^K \left\| \left([\chi(t)' \ -1] \begin{bmatrix} \omega_j - \omega_{s(t)} \\ \gamma_j - \gamma_{s(t)} \end{bmatrix} + 1 \right)_+ \right\|_2^2 + \beta_p \sum_{k=1}^K (\|\omega_k\|_2^2 + \gamma_k^2). \tag{13c}$$

Note that a regularization term is introduced in (13c) to make the problem strictly convex, with $\beta_p > 0$ being an additional tuning parameter.

	Description	Values
g	Gravitational acceleration	9.81 $\mathrm{m/s^2}$
m	Quarter-car mass	225 Kg
r	Wheel radius	30 cm
J	Wheel Inertia	1 $\mathrm{Nm^2}$
v	Longitudinal speed	50 m/s
F_z	Vertical force	2207.3 N

Table 1: Table of model parameters

By looking at the problem in (13), it can be seen that the cost is separable with respect to the parameters $\{\Theta_k\}_{k=1}^K$ and the ones of the PWA separator, once $\mathcal{S}$ is fixed. Based on this intuition, the design problem (13) is solved in an alternating fashion as summarized in Algorithm 1. Starting from an (eventually random) initial guess on the mode sequence $\mathcal{S}$, the approach consists of three main steps. Step 1.1 involves the design of the local governors. This is carried out by exploiting an instrumental variable scheme (see Söderström and Stoica (2002); Gilson and Van den Hof (2005) for more details), so as to cope with the non-white nature of the residual ε_u in (10). To this end, the loss in (13b) is slightly modified as

$$J_1^{IV}(\{\Theta_k\}_{k=1}^K, \mathcal{S}) = \left\| \sum_{t=1}^{T} z(t)\varepsilon_u(\Theta_{s(t)}, t) \right\|_2^2 + \lambda \sum_{k=1}^{K} \|\Theta_k\|_2^2, \tag{14}$$

where $z(t)$ is the *instrument* chosen by the user so as to be uncorrelated with the noise $v(t)$ acting on the measurements (see (2)). Step 1.2 can be efficiently handled by exploiting the Newton-like approach proposed in Breschi et al. (2016). Finally, step 1.3 involves the solution of an *unsupervised* clustering problem that can be tackled via dynamic programming tools, see Bertsekas (1999). These three steps are iterated until a user-defined stopping condition is met. We refer the reader to Breschi and Formentin (2020) for a more detailed description of the optimization approach underlying the above switching system design procedure.

4. The braking control case study

Let the data generating system (1) be described by the continuous-time *single corner model*:

$$\dot{\lambda}(t) = -\frac{1}{v}\left(\frac{1-\lambda(t)}{m} + \frac{r^2}{J}\right) F_z \mu(\lambda(t)) + \frac{r}{Jv} u(t), \tag{15}$$

where v [m/s] is the constant longitudinal speed, m [kg] is the quarter car mass, r [m] and J [Nm^2] are the wheel radius and moment of inertia, respectively. F_z [N] is the vertical force acting on the tire-road contact point. The values of variables in (15) are summarized in Table 1. We further impose the control variable $u(t)$ to saturate if either the minimum or the maximum allowed braking torque are reached, which are respectively equal to $\underline{u} = 0$ [Nm] of $\bar{u} = 2000$ [Nm]. The nonlinear relationship between between longitudinal friction coefficient μ and slip λ is described by the Burckhardt model (see Burckhardt (1993), Kiencke and Nielsen (2000b)), which has the following empirical expression:

$$\mu(\lambda, \alpha) = \alpha_1(1 - e^{\lambda\alpha_2}) + \lambda\alpha_3, \tag{16}$$

Road Condition	α_1	α_2	α_3
Asphalt dry	1.28	23.99	0.52

Table 2: Parameters of the Burckhardt model (16).

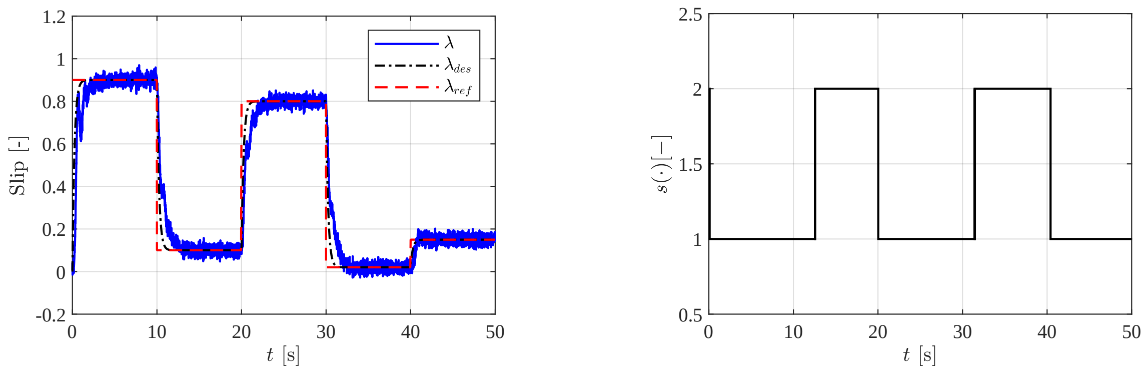

(*a*) Attained *vs* desired closed-loop behavior.

(*b*) Reference governor switching sequence.

Figure 4: Closed-loop performance: attained *vs* desired behavior for a piecewise constant reference.

where α is a constant parameter characterizing the road conditions. The values of the coefficients in α are defined in Table 2, and they correspond to dry asphalt conditions. This results in the same behavior depicted in Figure 1. We remark that neither the model describing the system dynamics in (15) nor the relationship between the friction coefficient μ and the slip λ in (16) are used when learning the reference governor.

The system in (15) is stabilized by the *proportional integral derivative* (PID) controller suggested in Savaresi and Tanelli (2010), namely

$$C(s) = 12000 \frac{\left(1 + \frac{1}{20}s\right)^2}{s\left(1 + \frac{1}{500}s\right)}. \tag{17}$$

This controller allows us to collect data and it represents the inner controller $\mathcal{C}$ in Figure 2, with its features assumed to be unknown when the reference governor is designed. The resulting closed loop system is fed by a piecewise constant slip set-point $\rho(t)$. Data are measured with a sampling time of $T_s = 0.01$ s, for a total of 100 s. Thus, the available dataset $\mathcal{D}_T = \{\rho(t), y(t)\}_{t=1}^{T}$ is constituted by $T = 10000$ samples. The output measurements are corrupted by a zero-mean additive white noise $v(t)$ with Gaussian distribution and standard deviation $\sigma = 0.015$, yielding a Signal-to-Noise Ratio (SNR)

$$\text{SNR} = 10 \log \left(\frac{\sum_{t=1}^{T} \lambda_o(t)^2}{\sum_{t=1}^{T} v(t)^2} \right) \approx 31.8 \text{ dB}. \tag{18}$$

The designed reference governor is

$$\mathcal{R}(\theta): \quad \rho(t) = \rho(t-1) + \theta_1^{s(t)} e(t) + \theta_2^{s(t)} e(t-1), \tag{19a}$$

so that it embeds an integral action to guarantee that constant references are tracked. The order of the reference governor in (19) has been chosen via cross-validation. The number of possible modes

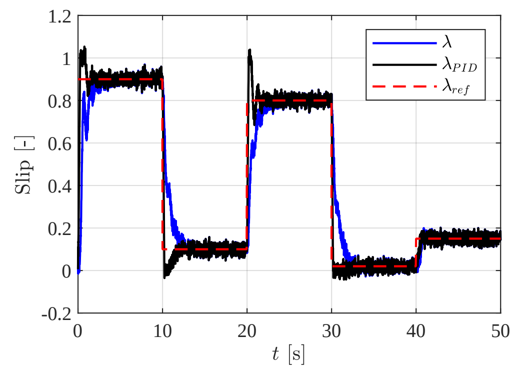

Figure 5: Closed-loop tracking performance: standalone PID *vs* nested scheme (PID plus switching reference governor).

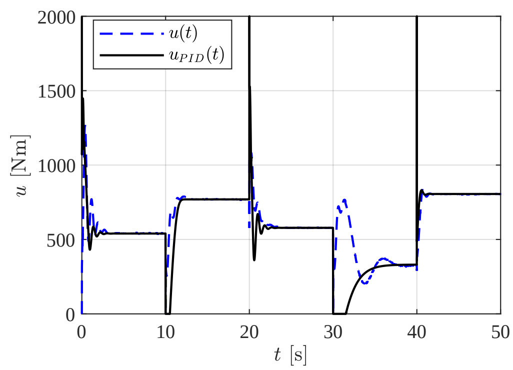

Figure 6: Braking torque (noiseless setting): standalone PID *vs* nested scheme (PID plus switching reference governor).

of the controller is set to $K = 2$, according to the number of actual operating conditions of the braking wheel system. We remark that this is the only prior on the system exploited in the design of $\mathcal{R}$. We impose the switching logic to be driven by vectors $\chi(t)$ of the form

$$\chi(t) = \begin{bmatrix}\rho(t-1) & y(t) & y(t-1)\end{bmatrix}'. \tag{19b}$$

The reference model $\mathcal{M}$ is the first order discrete time model:

$$\begin{cases} x_M(t+1) = 0.96 x_M(t) + r(t) \\ y_{des}(t) = 0.04 x_M(t), \end{cases} \tag{20}$$

which has a cutoff frequency of 0.8 Hz, allowing us to remove the undesired effects of the two resonance peaks at 5 Hz and 12 Hz, and contemporarily obtain a satisfactory behavior in either the stable and unstable region.

Algorithm 1 is executed for 20 randomly initialized switching sequences, for a maximum of $i_{max} = 100$ iterations for each initial condition. The controller is validated applying a piecewise constant set-point that explores both the stable and the unstable region over a horizon of 50 seconds. All the experiments are conducted within a noisy setting, with a noise equivalent to that of the training phase. The results displayed in Figure 4(*a*) show that the overall closed loop system is capable to properly track the reference, with the desired closed-loop behavior matched for both the stable and unstable mode. The occasional slowdown of the response reflects switches of the reference governor, that appropriately changes mode in correspondence of shifts between the stable and unstable region of the braking system. Nonetheless, the closed-loop system is able to recover the desired behavior reasonably fast.

We then compare the results obtained with the inner PID controller *alone* with the performance of the full nested strategy. The attained closed-loop slips are reported in Figure 5, which clearly shows that the data-driven reference governor *improves* the performance of the system. Indeed, it removes the undesired oscillatory behavior and overshoots characterizing the response obtained with the standalone PID (see Figure 5). At the same time, undesired saturations that may have a bad impact on the actuators are reduced, with a resulting smoother control action. Indeed, with the PID controller only, the input saturates 8% of the times (namely, 405 input samples out of 5000 are saturated). This percentage drops to 4% when the reference governor is used, i.e., 226 out of 5000 samples are saturated. Although most saturations are due to the effect of noise, there are exceptions linked to the additional control effort required by the standalone PID controller. This can be seen by looking at Figure 6, where we report the braking torques needed to follow the same piecewise constant reference of Figure 5 within a noiseless setting. It is clear that the use of the standalone PID implies additional control efforts for the closed-loop system to adapt to shifts in the operating region of the braking system due to changes of the set-point.

5. Conclusions

In this work, we propose a data-driven design procedure for switching reference governors that can be used to improve the closed-loop performance of an *a-priori* stabilized brake-by-wire system. By specifying the desired closed-loop behavior via a reference model, it is shown that the proposed approach allows us to track the desired closed-loop behavior and to decrease the overall control effort, while exploiting the full operational range of the actuators.

Future works will be devoted to assess the proposed approach within a more realistic experimental setup and to perform a sensitivity analysis with respect to the involved design parameters.

Acknowledgments

This project 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.

References

D. P. Bertsekas. *Nonlinear programming*. Athena scientific Belmont, 1999.

L. Breiman. Hinging hyperplanes for regression, classification, and functio approximation. *IEEE Transactions on Information Theory*, 39(3):999–1013, 1993.

V. Breschi and S. Formentin. Direct data-driven design of switching controllers. *International Journal of Robust and Nonlinear Control*, 30(15):6042–6072, 2020.

V. Breschi, D. Piga, and A. Bemporad. Piecewise affine regression via recursive multiple least squares and multicategory discrimination. *Automatica*, 73:155–162, 2016.

M. Burckhardt. *Fahrwerktechnik: Radschlupf-Regelsysteme: Reife- nverhalten, Zeitablaeufe, Messung des Drehzustands der Raeder, Anti-Blockier-System (ABS), Theorie Hydraulikkreislaeufe, Antriebs-Schlupf-Regelung (ASR), Theorie Hydraulikkreislaeufe, elektronische Regeleinheiten, Leistungsgrenzen, ausgefuehrte Anti-Blockier-Systeme und Antriebs-Schlupf-Regelungen*. Vogel, 1993.

M. C Campi, A. Lecchini, and S. M Savaresi. Virtual reference feedback tuning: a direct method for the design of feedback controllers. *Automatica*, 38(8):1337–1346, 2002.

C. Canudas-de-Wit, P. Tsiotras, E. Velenis, M. Basset, and G. Gissinger. Dynamic friction models for road/tire longitudinal interaction. *Vehicle System Dynamics*, 39(3):189–226, 2003.

A. Chakrabarty, K. Berntorp, and S. Di Cairano. Learning-based parameter-adaptive reference governors. In *2020 American Control Conference (ACC)*, pages 956–961, 2020.

S. Drakunov, U. Ozguner, P. Dix, and B. Ashrafi. Abs control using optimum search via sliding modes. *IEEE Transactions on Control Systems Technology*, 3(1):79–85, 1995.

M. Fliess and C. Join. Model-free control and intelligent PID controllers: Towards a possible trivialization of nonlinear control? *IFAC Proceedings Volumes*, 42(10):1531 – 1550, 2009. 15th IFAC Symposium on System Identification.

S. Formentin, P. De Filippi, M. Corno, M. Tanelli, and S.M Savaresi. Data-driven design of braking control systems. *IEEE Transactions on Control Systems Technology*, 21(1):186–193, 2011.

S. Formentin, C. Novara, S.M. Savaresi, and M. Milanese. Active braking control system design: The D2-IBC approach. *IEEE/ASME Transactions on Mechatronics*, 20(4):1573–1584, 2015.

S. Formentin, D. Piga, R. Tóth, and S. M Savaresi. Direct learning of LPV controllers from data. *Automatica*, 65:98–110, 2016.

S. Formentin, M.C. Campi, A. Carè, and S.M. Savaresi. Deterministic continuous-time virtual reference feedback tuning (VRFT) with application to PID design. *Systems & Control Letters*, 127:25–34, 2019.

M. Gilson and P. Van den Hof. Instrumental variable methods for closed-loop system identification. *Automatica*, 41(2):241 – 249, 2005. ISSN 0005-1098.

U. Kiencke and L. Nielsen. *Automotive Control Systems: For Engine, Driveline and Vehicle*. Springer-Verlag, Berlin, Heidelberg, 1st edition, 2000a. ISBN 3540669221.

U. Kiencke and L. Nielsen. Automotive control systems, SAE international, 2000b.

K. Liu, N. Li, D. Rizzo, E. Garone, I. Kolmanovsky, and A. Girard. Model-free learning to avoid constraint violations: An explicit reference governor approach. In *2019 American Control Conference (ACC)*, pages 934–940, 2019.

J. Martinez and C. Canudas-de-Wit. A safe longitudinal control for adaptive cruise control and stop-and-go scenarios. *IEEE Transactions on Control Systems Technology*, 15(2):246–258, 2007. doi: 10.1109/TCST.2006.886432.

M. Radac, R. Precup, and R. Roman. Anti-lock braking systems data-driven control using Q-learning. In *2017 IEEE 26th International Symposium on Industrial Electronics (ISIE)*, pages 418–423, 2017.

S.M. Savaresi and M. Tanelli. *Active Braking Control Systems Design for Vehicles*. Advances in Industrial Control. Springer London, 2010. ISBN 9781849963503.

T. Söderström and P. Stoica. Instrumental variable methods for system identification. *Circuits, Systems and Signal Processing*, 21(1):1–9, 2002.

C.J. C.H. Watkins and P. Dayan. Q-learning. *Machine Learning*, 8(3):279 – 292, 1992.

P. Wellstead and N. Pettit. Analysis and redesign of an antilock brake system controller. *IEE Proceedings-Control Theory and Applications*, 144(5):413–426, 1997.

J. Yi, L. Alvarez, R. Horowitz, and C. Canudas de Wit . Adaptive emergency braking control using a dynamic tire/road friction model. In *Proceedings of the 39th IEEE Conference on Decision and Control (Cat. No.00CH37187)*, volume 1, pages 456–461 vol.1, 2000. doi: 10.1109/CDC.2000.912806.