Proceedings of Machine Learning Research vol 120:1-9, 2020

Direct data-driven control with embedded anti-windup compensation

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

Input saturation is an ubiquitous nonlinearity in control systems and arises from the fact that all actuators are subject to a maximum power, thereby resulting in a hard limitation on the allowable magnitude of the input effort. In the scientific literature, anti-windup augmentation has been proposed to recover the desired linear closed-loop dynamics during transients, but the effectiveness of such a compensation is strongly linked to the accuracy of the mathematical model of the plant. In this work, it is shown that a feedback controller with embedded anti-windup compensator can be directly identified from data, by suitably extending the existing data-driven design theory. The effectiveness of the resulting method is illustrated on a benchmark simulation example.

Keywords: direct data-driven control, anti-windup, input saturation

1. Introduction

In real-world control systems, the power of actuators is always limited, thus not allowing the required control inputs to be always fully applicable to the plant. While most of the existing industrial applications address saturation problems by ensuring that the actuators are powerful enough not to reach the saturation limits during operations, the demand for increasingly performing control solutions is calling upon design techniques that explicitly account for these input limitations.

In traditional industrial and process controls, a small-signal controller is first designed disregarding input saturation based on a desired reference behavior and, in a second step, an *anti-windup compensator* is designed so to guarantee the local preservation of the small signal response as well as the asymptotic recovery of performance whenever the saturation limits are reached during the transients. Anti-windup design is especially desirable when dealing with linear plants with input saturation, because in that case small-signal design can be performed by focusing on a linear plant thereby using many well established design techniques. Then, the nonlinearity arising from input saturation is addressed in the anti-windup design phase, see, e.g., Tarbouriech et al. (2011); Galeani et al. (2009) and the references therein. As clearly reported in Zaccarian and Teel (2011), all the major design approaches to anti-windup design not only require the knowledge of the saturation limits, but they also rely on the hypothesis that an accurate mathematical model of the system is available. While the former assumption is practically always verified in most control applications, the latter might be too strict in case some prior physical knowledge is not given and the model must be identified from data.

An alternative perspective to the design problem that could be suited in the above situation is *direct data-driven control* (see Hou and Wang (2013); Formentin et al. (2014)), where data collected on the plant to control are directly mapped onto the controller parameters, without first identifying a model of the system. To the best of the authors' knowledge, the existing techniques are suited for the off-line design of the feedback controls (e.g., for reference tracking or disturbance rejection, Formentin and Karimi (2012); Markovsky (2016)) but there is no way to embed an anti-windup compensator so as to explicitly deal with input saturation.

In this paper, the direct data-driven approach for linear control design introduced in Formentin et al. (2016) is extended to allow for the joint design of the feedback controller and the anti-windup compensator. The resulting procedure turns out to yield the desired closed-loop performance, in case the reference model describing the specifications is achievable using a controller within the given parametric class. It should be mentioned here that a data-driven approach to deal with constrained signals has been proposed in Piga et al. (2017). However, the method presented here is substantially different from the one presented in Piga et al. (2017), since a 2 degrees of freedom architecture is therein needed with a real-time optimization task running as a reference governor. Instead, in this work, a simple parametric controller is used, whose parameters can be computed off-line from data via simple regularized least-squares.

The remainder of the paper is organized as follows. The problem of data-driven control design in presence of actuator saturation is formalized in Section 2, while the proposed solution is discussed in Section 3. Section 4 reports the results of a numerical case study, showing the effectiveness of the proposed strategy. Conclusions and directions for future research are finally indicated in Section 5.

2. Problem setting

Let P be a *single-input single-output* (SISO) system, with dynamics described by the following (unknown) *input-output* (IO) relation

$$P: \quad A(q^{-1})y_{\mathrm{o}}(t) = B(q^{-1})g(u(t)), \tag{1a}$$

where $u(t) \in \mathbb{R}$ and $y_{\mathrm{o}}(t) \in \mathbb{R}$ are the input and the *noiseless* output to P at time t. The polynomials $A(q^{-1})$ and $B(q^{-1})$ in the shift operator q^{-1} (*i.e.,* $q^{-1}u(t) = u(t-1)$), are defined as

$$A(q^{-1}) = 1 + \sum_{i=1}^{n_a} a_i q^{-i}, \quad B(q^{-1}) = \sum_{j=0}^{n_b} b_j q^{-j}, \tag{1b}$$

with orders $n_a, n_b \in \mathbb{N}$ indicating the dynamical order of the system, and $g : \mathbb{R} \to \mathbb{R}$ being the nonlinear function

$$g(u) = \begin{cases} -\alpha & \text{if } u \leq -\alpha, \\ u & \text{if } -\alpha < u < \alpha, \\ \alpha & \text{if } u \geq \alpha, \end{cases} \tag{1c}$$

with $\alpha \in \mathbb{R}_{>0}$. Apart from the above constant α, the plant dynamics is supposed to be fully unknown. However, we assume that an experiment can be performed by feeding the system with a bounded input sequence $\{u(t)\}_{t=1}^{T}$ and that we can measure the corresponding *noisy* output

$$y(t) = y_{\mathrm{o}}(t) + v(t), \;\; t = 1, \ldots, T, \tag{1d}$$

with $v(t) \in \mathbb{R}$ being the realization at time t of a zero-mean, white noise sequence. We further suppose the system to be *bounded-input bounded-output* (BIBO) stable, namely

$$\sup_{t \geq 0} |u(t)| < +\infty \Rightarrow \sup_{t \geq 0} |y_{\mathrm{o}}(t)| < +\infty.$$

In this work, we aim at learning a controller for the system *directly from input/output data*, rather than identifying P and then solving a model-based control problem. The controller is designed to

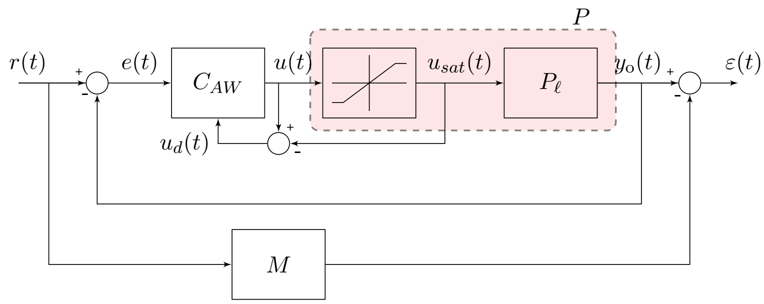

Figure 1: The proposed closed-loop matching scheme. The actual plant P is highlighted in red.

match a *linear time invariant* (LTI) reference model $M(q^{-1})$, which describes the desired closed-loop behavior. Nonetheless, differently from existing data-driven techniques, we want to exploit our insights on the nonlinear function $g(\cdot)$ in (1c) to embed an *anti-windup* action in the controller C_{AW}. Therefore, we design C_{AW} with a (linear) structure inspired by state-of-the-art anti-windup schemes of Zaccarian and Teel (2011), namely

$$C_{AW}: \quad A_c(q^{-1})u(t) = B_c(q^{-1})e(t) + B_c^{AW}(q^{-1})u_d(t), \tag{2a}$$

where $u_d(t)$ is the difference between the input $u(t)$ fed to the system and its saturated counterpart $u_{sat}(t)$, and the polynomials $A_c(q^{-1})$, $B_c(q^{-1})$ and $B_c^{AW}(q^{-1})$ are chosen as:

$$A_c(q^{-1}) = 1 + \sum_{i=1}^{n_a^c} a_i^c q^{-i}, \quad B_c(q^{-1}) = \sum_{j=0}^{n_b^c} b_j^c q^{-j}, \quad B_c^{AW}(q^{-1}) = \sum_{k=1}^{n_b^{AW}} b_k^{AW} q^{-k}. \tag{2b}$$

Let $\theta \in \mathbb{R}^n$ be the vector of the unknown controller parameters, *i.e.,*

$$\theta = \begin{bmatrix} -a_1^c & \cdots - a_{n_a^c}^c & b_0^c & b_1^c & \cdots & b_{n_b^c}^c & b_1^{AW} & \cdots & b_{n_b^{AW}}^{AW} \end{bmatrix}', \tag{3}$$

with $n = n_a^c + n_b^c + n_b^{AW}$, and let the regressor $x(t) \in \mathbb{R}^n$ be

$$x(t) = \begin{bmatrix} u(t-1) & \cdots & u(t-n_a^c) & e(t) & e(t-1) & \cdots & e(t-n_b^c) & u_d(t-1) & \cdots & u_d(t-n_b^{AW}) \end{bmatrix}', \tag{4}$$

so that the controller dynamics in (2) can be equivalently defined as

$$u(t) = x(t)'\theta. \tag{5}$$

According to the closed-loop matching scheme reported in Figure 1, our ultimate design goal reduces to the minimization of the *mismatch error*

$$\varepsilon(t) = y_o(t) - M(q^{-1})r(t), \tag{6}$$

via the solution of the following optimization problem

$$\min_{\theta,\varepsilon} \ \|\varepsilon\|_{\ell_2}^2 \tag{7a}$$

$$\text{s.t. } \varepsilon(t) = y_o(t) - M(q^{-1})r(t), \ \ t = 1, \ldots, T, \tag{7b}$$

$$A(q^{-1})y_o(t) = B(q^{-1})g(u(t)), \ \ t = 1, \ldots, T, \tag{7c}$$

$$u(t) = x(t)'\theta, \ \ t = 1, \ldots, T. \tag{7d}$$

3. Direct data-driven controller design with embedded anti-windup

The minimization problem in (7) cannot be readily solved, due to the critical dependences on: (i) the *specific* reference to be tracked (see (7b) and (7d)); (ii) the *unknown* plant dynamics (as in (7b)). Not to fix the reference to be tracked a priori (with the risk of over-fitting the controller to a specific scenario), we exploit the reference model $M(q^{-1})$ to annihilate the dependence on $r(t)$ via the construction of a *fictitious reference*

$$r_f(t) = M^{\dagger}(y_{\rm o}(t) - \varepsilon(t)), \tag{8}$$

with $M^{\dagger}$ denoting the left inverse of the reference model $M(q^{-1})$, *i.e.,* $M^{\dagger}M = 1$. The new reference r_f depends on the noiseless outputs $y_{\rm o}(t)$ and the optimization variables $\varepsilon(t)$. Therefore, it does not require one to fix the reference to be tracked beforehand, ultimately leading to a more *general* controller. On the other hand, it requires the computation of the left inverse $M^{\dagger}$, which might be not a trivial operation especially when considering nonlinear reference models (see Formentin et al. (2016) for further details). By replacing $r(t)$ with the fictitious signal $r_f(t)$ and defining $\xi_{\rm o} = (M^{\dagger} - 1)y_{\rm o}$, the constraint in (7d) is modified as

$$u(t) = x_{f,{\rm o}}(t)'\theta - \tilde{x}_f(t)'\tilde{\theta}, \tag{9a}$$

where

$$x_{f,{\rm o}}(t) = \big[u(t-1) \cdots u(t-n_a^c) \;\; \xi_{\rm o}(t) \;\; \xi_{\rm o}(t-1) \cdots \xi_{\rm o}(t-n_b^c) \;\; u_d(t-1) \cdots u_d(t-n_b^{AW})\big]', \tag{9b}$$

$$\tilde{\theta} = \begin{bmatrix} b_0^c & b_1^c & \cdots & b_{n_b^c}^c \end{bmatrix}', \tag{9c}$$

$$\tilde{x}_f(t) = \big[\varepsilon(t-1) \;\; \cdots \;\; \varepsilon(t-n_b^c)\big]'. \tag{9d}$$

While the reference model and the definition of the mismatch error ε are exploited to eliminate the dependence on a prefixed reference signal, the constraint on the plant dynamics is dropped by replacing the inputs $u(t)$ and noiseless outputs $y_{\rm o}(t)$ in (7) with the available data $\mathcal{D}_T = \{u(t), y(t)\}_{t=1}^T$. The resulting problem is given by:

$$\min_{\theta,\varepsilon} \; \|\varepsilon\|_{\ell_2}^2 \tag{10a}$$

$$\text{s.t. } \tilde{x}_f(t)'\tilde{\theta} = x_f(t)'\theta - u(t), \;\; t = 1, \dots, T, \tag{10b}$$

where the remaining constraint is obtained via a straightforward manipulation of (9a), and $x_{f,{\rm o}}(t)$ is replaced by $x_f(t)$, defined as in (9b) with $\xi_{\rm o}$ substituted with $\xi = (M^{\dagger} - 1)y$. Since the measured outputs are noisy, this problem is indeed an approximation of the original one in (7).

Note that the left-hand side of (10b) is bi-convex in the unknowns, as it is the product of a subset of the controller parameters and a stack of delayed values of the mismatch error. This issue is bypassed via the optimization of the squared norm of this bi-convex term instead of the minimization of $\|\varepsilon\|_{\ell}^2$, namely by solving the following regularized least-squares problem

$$\min_{\theta} \; \frac{1}{T} \sum_{t=1}^{T} (x_f(t)'\theta - u(t))^2 + \lambda \|\theta\|_{\ell_2}^2. \tag{11}$$

This problem is solved through an *instrumental variable* (IV) scheme Gilson and Van den Hof (2005); Söderström and Stoica (1988); Ljung (1986), in order to avoid the trivial solution $\tilde{\theta} = 0$, while accounting for the noise affecting the available data. Notice that this requires the proper selection of an instrument $\zeta \in \mathbb{R}^n$. According to standard instrumental variable schemes, ζ can be

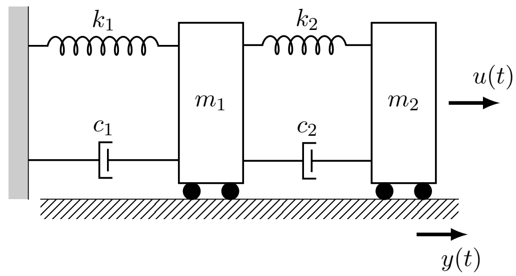

Figure 2: Scheme of the physical system to be controlled.

Table 1: Parameters of the system P.

m_1 [kg]	m_2 [kg]	c_1 [N/m^2]	c_2 [N/m^2]	k_1 [N/m]	k_2 [N/m]
1	0.5	0.2	0.5	1	0.5

defined as a new instance of the regressor x_f, obtained by either running a new experiment on the plant with the input sequence $\{u(t)\}_{t=1}^{T}$ (thus considering a different realization of the noise) or by identifying an approximate (LTI) model for P and then simulating the model with the given input $u(t)$. Once $\{\zeta(t)\}_{t=1}^{T}$ is retrieved, it is easy to prove that the estimated parameters are given by

$$\theta_{IV} = \left[\sum_{t=1}^{T} \zeta(t) x_f(t)' + \lambda I\right]^{-1} \left(\sum_{t=1}^{T} \zeta(t) u(t)\right). \tag{12}$$

We stress here that the regularization term $\lambda \in \mathbb{R}_{\geq 0}$ in (11) is introduced to improve the statistical properties of the learned controller, while enforcing the parameters not to be excessively large.

4. A numerical case study

To assess the performance of the proposed data-driven anti-windup control design, we consider the system depicted in Figure 2 taken from Caré et al. (2019), whose *continuous* dynamical model is given by the transfer function in the Laplace variable $s \in \mathbb{C}$:

$$P(s) = \frac{m_1 s^2 + (c_1 + c_2)s + (k_1 + k_2)}{(m_1 s^2 + (c_1 + c_2)s + k_1 + k_2)(m_2 s^2 + c_2 s + k_2) - (c_2 s + k_2)^2}, \tag{13}$$

with parameters reported in Table 1. The transfer function $P(s)$ describes the relationship between the force u that can be applied to the system to pull the masses away from the wall, and the position y of the 2nd mass m_2. We feed the system with an input sequence of length 5000 samples given by the sum of a uniformly distributed random sequence in the interval $[-1,\ 1]$ and a square wave with amplitude 1.2 and period of approximately 12 seconds. The measured output is corrupted by an additive, zero-mean Gaussian white noise with variance equal to 0.06. The effect of the noise on the data is assessed through the *Signal-to-Noise Ratio* (SNR)

$$\text{SNR} = \frac{\sum_{t=1}^{T} y_o(t)^2}{\sum_{t=1}^{T} v(t)^2} \approx 20 \text{ dB}. \tag{14}$$

Suppose that we aim the closed-loop system to behave as the following first order reference model

$$M(s) = \frac{1}{5s + 1}, \tag{15}$$

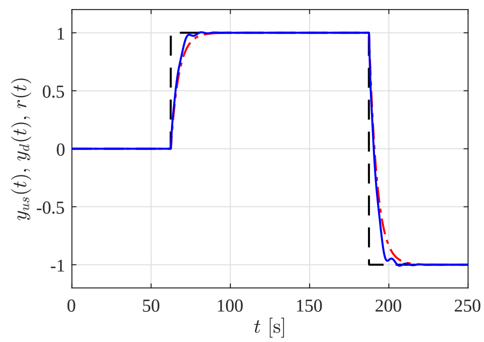

Figure 3: Benchmark closed-loop behavior. Reference r (black dashed), desired response y_d (dotted dashed red) and attained closed-loop output y_{uc} (blue) without actuator saturation.

and let us initially assume that $g(u(t)) = u(t)$ for all $t \in \mathbb{N}$, so that the plant P actually corresponds to P_ℓ and, thus, it is linear. Under this hypothesis, we may exploit state-of-the-art data-driven model reference techniques like the Virtual Reference Feedback Tuning approach (VRFT, Formentin et al. (2019)) to train a *Proportional-Integral-Derivative* (PID) controller, namely, in discrete time:

$$C_{PID}(q^{-1}) = K_p + K_i \frac{T_s}{2} \frac{1+q^{-1}}{1-q^{-1}} + K_d \frac{2}{T_s} \frac{1-q^{-1}}{3-q^{-1}}, \tag{16}$$

where $T_s = 0.1$ s is the system sampling time. The (noiseless) closed-loop output obtained when tracking a square wave is reported in Figure 3, along with the desired closed-loop response. Given the fairly good results attained in this setting, in the rest of the section these results are considered as a benchmark for performance assessment.

Suppose now that the actuator saturates once the input provided by the controller exceeds the threshold $\alpha = 0.5$, so that P is described by the model in (1). We stress that in our framework the saturation is supposed to be *known* beforehand, so it was already exploited in selecting the input sequence used for data generation to acquire as much information as possible on all the operating conditions of P.

In this setting, we train a controller C_{AW} with integral action

$$u(t) = -\sum_{i=1}^{2} a_i^c u(t-i) + \sum_{j=0}^{2} b_j^c e_{int}(t-j) + \sum_{k=1}^{2} b_k^{AW} u_d(t-k), \tag{17a}$$

$$e_{int}(t) = e_{int}(t-1) + (r(t) - y(t)), \tag{17b}$$

so to guarantee zero steady-state error. When learning the controller, we generate the instrument by identifying a simple 4-th order LTI model of P and simulating its response to the chosen input. The parameters of (17) are then retrieved by solving (12), with $\lambda = 0.25$ (found via cross-validation). Along with the benchmark unsaturated case, the performance attained with the proposed anti-windup controller is compared with the one obtained with two different PID controllers $C_{1,PID}$ and $C_{2,PID}$ learned through VRFT Formentin et al. (2019). In the first case, we exploit the original inputs $\{u(t)\}_{t=1}^{T}$ only, thus completely neglecting our insights on the saturation. In the second case, we use saturated inputs $\{u_{sat}(t)\}_{t=1}^{T}$ only, not considering the chosen (informative) input sequence. Since we want to concurrently design the tracking control and the anti-windup action, we do not enforce the PID structure on (17). Nonetheless, for the comparison to be fair

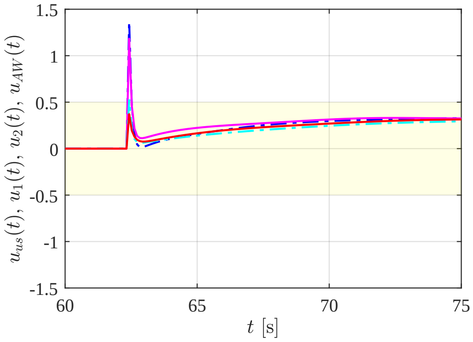

(a) Detail of the input.

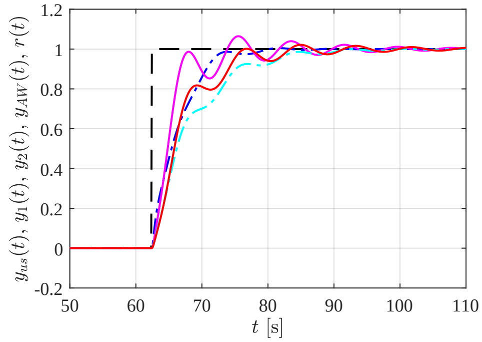

(b) Detail of the reference (dashed black) vs closed-loop output.

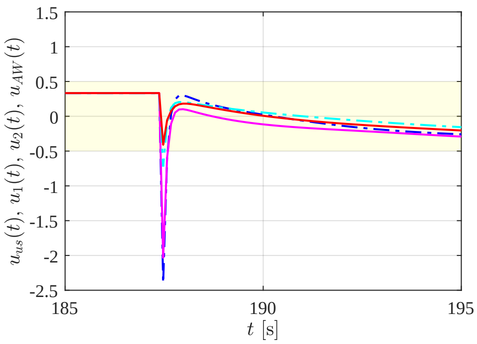

(c) Detail of the input.

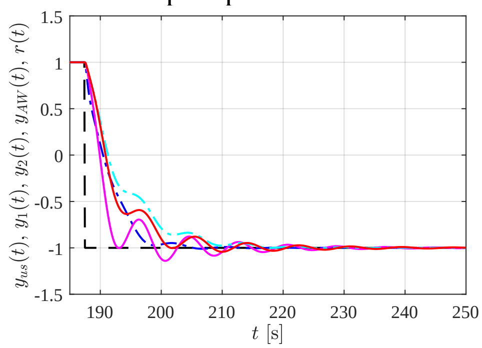

(d) Detail of reference (dashed black) vs closed-loop output.

Figure 4: Control input [left panels] and output sequences [right panels]. The illustrated signals are as follows: unsaturated benchmark, u_{us} and y_{us} (dotted dashed blue); input and output with $C_{1,PID}$ (magenta); input and output with $C_{2,PID}$ (dotted dashed cyan); input and output with C_{AW} (red). The yellow regions indicate where the inputs are not saturated.

we always generate the instrument as explained previously and, in learning (17), we preprocess all data with the optimal filter exploited when training $C_{1,PID}$[1]. The closed-loop inputs and (noiseless) outputs obtained with the different controllers when tracking a square wave are shown in Figure 4. As expected, when the saturated input is not used in the design phase, the control input visibly exceeds the saturation thresholds (see Figures 4(a)-4(c)). At the same time, by solely exploiting the saturated inputs, the control action might be too conservative, while still slightly exceeding the thresholds. On the other hand, despite the differences in the controller structures, the use of the proposed anti-windup like strategy leads to a control input that resembles the one obtained in the unsaturated benchmark, while never violating the saturation limits (see Figure 4(c)). This behavior affects the final tracking performance attained in closed-loop. Indeed, as visible in Figure 4(b) and Figure 4(d), the output obtained when exploiting $C_{1,PID}$ tends to visibly oscillate before converging to the benchmark behavior, since sudden changes in the reference usually lead to input saturations. On the other hand, the use of $C_{2,PID}$ allows the input to recover from saturation quite rapidly at

1. This choice is driven by the fact that both C_{AW} and C_1 are learned by fitting the unsaturated input sequence $\{u(t)\}_{t=1}^{T}$.

Table 2: Obtained performance indexes.

	RMSE_r	RMSE_d	Saturation violations
$C_{1,PID}$	0.19	0.10	Yes
$C_{2,PID}$	0.25	0.04	Yes
C_{AW}	0.23	0.03	No

the price of a slower convergence to the set point. Instead, the output obtained by closing the loop with C_{AW} adheres the most to the one attained when removing the saturation, thus proving the benefits of the proposed approach. These results are confirmed by the performance indexes reported in Table 2, namely

$$\text{RMSE}_r = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(r(t)-y(t))^2}, \quad \text{RMSE}_d = \sqrt{\frac{1}{T}\sum_{t=1}^{T}(y_d(t)-y(t))^2}, \tag{18}$$

which provide quantitative information on the capabilities of the closed-loop system to track the considered reference signal and the behavior of the reference model, respectively.

5. Conclusions

In this paper, we have proposed a data-driven control design method suited to handle actuator saturation. The proposed dynamic controller is trained by solving a tailored model-reference design problem and it combines past tracking errors and differences between nominal and saturated inputs to generate the next control action. The numerical results show that the proposed controller allows to attain a trade-off between tracking performance and saturation violation, with the resulting control action being not excessively aggressive nor too conservative.

Given the importance of a proper selection of the input used for data generation, future work will be focused on how to properly design the controller identification experiment.

References

A. Caré, F. Torricelli, M.C. Campi, and S.M. Savaresi. A toolbox for virtual reference feedback tuning (VRFT). In *2019 18th European Control Conference (ECC)*, pages 4252–4257, June 2019.

S. Formentin and A. Karimi. A data-driven approach to mixed-sensitivity control with application to an active suspension system. *IEEE Transactions on Industrial Informatics*, 9(4):2293–2300, 2012.

S. Formentin, K. van Heusden, and A. Karimi. A comparison of model-based and data-driven controller tuning. *International Journal of Adaptive Control and Signal Processing*, 28(10):882–897, 2014.

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.

S. Galeani, S. Tarbouriech, M. Turner, and L. Zaccarian. A tutorial on modern anti-windup design. *European Journal of Control*, 15(3):418 – 440, 2009.

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

Zhong-Sheng Hou and Zhuo Wang. From model-based control to data-driven control: Survey, classification and perspective. *Information Sciences*, 235:3–35, 2013.

L. Ljung. *System Identification: Theory for the User*. Prentice-Hall, Inc., USA, 1986.

I. Markovsky. A missing data approach to data-driven filtering and control. *IEEE Transactions on Automatic Control*, 62(4):1972–1978, 2016.

D. Piga, S. Formentin, and A. Bemporad. Direct data-driven control of constrained systems. *IEEE Transactions on Control Systems Technology*, 26(4):1422–1429, 2017.

T. Söderström and P. Stoica. *System Identification*. Prentice-Hall, Inc., USA, 1988.

S. Tarbouriech, G. Garcia, J.M. Gomes da Silva Jr., and I. Queinnec. *Stability and stabilization of linear systems with saturating actuators*. Springer-Verlag London Ltd., 2011.

L. Zaccarian and A.R. Teel. *Modern Anti-windup Synthesis: Control Augmentation for Actuator Saturation.* Princeton University Press, 2011.