doi: 10.82124/CEAS-GNC-2026-083 CEAS-GNC-2026-083

EuroGNC Madrid, Spain May 5^{th}-7^{th} 2026

Attitude control design for a multirotor UAV: a predictor-based approach

Marco Lovera Professor, Department of Aerospace Science and Technology, Politecnico di Milano, Milano, Italy. marco.lovera@polimi.it

ABSTRACT

In this paper, the problem of designing the attitude control law of a multirotor Unmanned Aerial Vehicle (UAV) is considered and a data-driven approach is proposed. With respect to previous work, the data-driven approach considered in this paper is the novel Virtual Reference Predictor-Based (VRPB) method, which will be used to design the controller based on data collected during closed-loop operation. Simulation results generated using a model of the attitude dynamics of a quadrotor UAV are used to illustrate the performance of the proposed design method.

Keywords: UAV, Attitude control, Data-driven control

1 Introduction

Small-scale Unmanned Aerial Vehicles (UAVs), and in particular multirotor ones, have been studied extensively in view of the great potential for a large number of applications. For most problems of practical interest, requirements in terms of pointing and positioning performance require a careful design of the control laws. While nonlinear control design approaches have been considered in the literature (see, *e.g.*, [1, 2] for recent survey papers), for civil applications such as surveillance, mapping, video and photography linear controllers are usually adopted. As in these settings hover and near-hover operations are representative conditions, cascaded PID laws are usually employed for attitude control thanks to their widely proven reliability and ease of implementation. As far as controller tuning is concerned, however, model-based methods suffer from the fact that the mathematical modelling of quadrotors is particularly challenging due to the non-trivial characterization of the aerodynamics and of the actuators and sensors dynamics (see [3]). For this reason data-driven tuning methods, which have been developed in the last two decades in the control community, offer an interesting alternative. These control design tools are especially appealing when *a priori* knowledge about the plant model is limited, when an accurate modeling of the system is too expensive or when fast deployment of the control system is an important requirement, since they allow the direct tuning of the controller parameters from experimental input-output data. Among the data-driven methods available in the literature, a coarse classification can be made between iterative (*e.g.*, the Iterative Feedback Tuning (IFT) [4]) and single-shot (non-iterative) methods (*e.g.*, the Virtual Reference Feedback Tuning (VRFT) [5], the Correlation-Based Tuning (CbT) [6, 7]). Recent advances on the VRFT method, which is the approach adopted in this work, can be found *e.g.*, in [8–10], while application studies are available, *e.g.*, in [11, 12]. Recently (see [13] and references therein) the VRFT algorithm has been considered to tune the attitude controller parameters of quadrotors; in particular, an extension of VRFT allowing the direct tuning of a cascade controller configuration with a single set of input-output data has been succesfully applied, both for SISO and MIMO problems.

Recent developments in virtual reference data-driven methods, however, are exploring the integration of more advanced model identification methods within the data-driven framework, such as the Prediction-Based Subspace Identification (PBSID) method. PBSID belongs to the family of subspace identification methods and has been developed specifically to guarantee closed-loop consistency by relying on finite-horizon predictors that incorporate future inputs [14, 15]. By subtracting the estimated effect of these future inputs and recovering states through a singular-value decomposition, PBSID yields asymptotically unbiased estimates of state-space models from closed-loop data. Several variants exist, such as the VARX-based formulation and the VARMAX-based PBSIDopt, which improves finite-sample properties by avoiding excessively long horizons [16, 17]. Comprehensive surveys of predictor-based methods highlight that PBSID and its relatives are asymptotically equivalent to appropriately weighted prediction-error methods while retaining the computational efficiency of subspace algorithms [18]. Embedding PBSID within VRFT thus provides a regression step that (i) enables unstructured state-space controllers, (ii) improves statistical reliability with closed-loop data, and (iii) preserves VRFT's one-shot tuning philosophy. In view of this, in this paper a closed-loop approach to data-driven design of the attitude control laws for a multirotor UAV is presented. With respect to previous work, the Virtual Reference Predictor-Based (VRPB) framework is adopted and the achievable performance is illustrated by means of simulation results obtained for a small-scale quadrotor.

The paper is organised as follows: Section 2 provides some background on the VRFT method and on the PBSID model identification algorithm and introduces VRPB. Subsequently, Section 3 presents simulation results while Section 4 concludes the paper.

2 Data-driven control law design

2.1 Introduction to Virtual Reference Feedback Tuning

Consider a linear time-invariant discrete-time system $P(z)$, where z denotes the forward time-shift unit (*i.e.*, $zx(t) = x(t+1)$), a parametrized controller class $C(\theta) = \{C(z,\theta), \theta \in \mathbb{R}^n\}$, and a given target closed-loop behaviour $M(z)$. The control aim of data-driven methods is the minimization of the weighted $\mathcal{L}_2$-norm of the mismatch between $M(z)$ and the actual closed-loop system:

$$J_{MR}(\theta) = \left\| \left(\frac{P(z)C(z,\theta)}{1 + P(z)C(z,\theta)} - M(z) \right) W(z) \right\|_2^2, \tag{1}$$

where $W(z)$ is a weighting function chosen by the user. In data-driven approaches the model-reference problem (1) is solved using only a set of available measurements $d_N = \{u(t), y(t)\}_{t=1,\ldots,N}$, where N is the length of the data-set and u, y are the input and output to the plant $P(z)$, respectively.

Consider the reference signal $r(t)$ that would feed the system in closed-loop operation when the closed-loop model is $M(z)$ and the output is the measured $y(t)$. Such a signal is called *virtual reference* and is such that $y(t) = M(z)r(t)$. A good controller (making the closed-loop as close as possible to $M(z)$) is then the one that produces the input sequence of the experiment $u(t)$ when it is fed by the error signal $e(t) = r(t) - y(t)$.

Formally, the cost criterion minimized by the VRFT algorithm is the following:

$$J_{VR}^N(\theta) = \frac{1}{N} \sum_{t=1}^{N} \left(u_L(t) - C(z,\theta) e_L(t) \right)^2, \tag{2}$$

where $u_L(t)$ and $e_L(t)$ are suitably filtered versions of $u(t)$ and $e(t)$. The filter $L(z)$ is chosen such that the cost function (2) is a local approximation of the criterion (1) in the neighborhood of the minimum point [5].

VRFT has been extended to deal with multiple nested loops architectures in [19, 20] and to MIMO systems. While VRFT provides a direct approach to data-driven controller design, in its classical form it relies on a regression in which the regressor, the virtual error, is computed from the measured output, and is therefore affected by measurement noise. This leads to an inherent errors-in-variables (EIV) problem that can bias the resulting controller estimate and therefore requires a careful treatment (see the cited references for details). Moreover, when data are collected under feedback or when a MIMO problem has to be solved, the closed-loop correlations and the issue of specifying the structure of the controller can further affect the quality of the result.

These limitations motivate the introduction of a more general identification approach within the VRFT framework, capable of handling unstructured models and closed-loop data in a simple way while retaining the method's one-shot nature.

2.2 Predictor-based subspace identification

Subspace identification methods are a class of model identificaton techniques which leverage the properties of linear systems and tackle the identification problem in terms of algebraic equations constructed from measured data. These methods are well suited for MIMO systems, do not require an initial guess for the model nor a parametrization of its matrices, and have strong consistency properties regarding the effect of noise.

For DT systems, the Predictor-Based Subspace Identification (PBSID, see the references in the Introduction) method is a well known approach that provides excellent results in practical scenarios. A short summary of the main step of the PBSID method is provided in the following, starting from the definition of the considered model class which is the one of LTI systems in innovation form

$$\begin{aligned} \hat{x}(t+1) &= A\hat{x}(t) + Bu(t) + Ke(t) \\ y(t) &= C\hat{x}(t) + Du(t) + e(t). \end{aligned} \tag{3}$$

where K is the steady-state Kalman gain and the innovation $e(t)$ is defined as

$$e(t) = y(t) - C\hat{x}(t) - Du(t). \tag{4}$$

Substituting the measurement equation in the state equation leads to

$$\hat{x}(t+1) = A\hat{x}(t) + Bu(t) + Ky(t) - KC\hat{x}(t) - KDu(t) = \bar{A}\hat{x}(t) + \bar{B}\begin{pmatrix} u(t) \\ y(t) \end{pmatrix}, \tag{5}$$

with $\bar{A} = A - KC$ and $\bar{B} = \begin{bmatrix} B - KD & K \end{bmatrix}$. Let the extended measurement $z(t) = \begin{pmatrix} u(t) & y(t) \end{pmatrix}^T$. Multiple-step ahead predictions of the state can be computed as

$$\hat{x}(t+2) = \bar{A}x(t+1) + \bar{B}z(t+1) = \bar{A}^2 x_k + \bar{A}\bar{B}z(t) + \bar{B}z(t+1). \tag{6}$$

Hence, in general,

$$\hat{x}(t+p) = \bar{A}^p \hat{x}(t) + \begin{bmatrix} \bar{B} & \bar{A}\bar{B} & \cdots & \bar{A}^{p-1}\bar{B} \end{bmatrix} \begin{pmatrix} z(t+p-1) \\ z(t+p-2) \\ \vdots \\ z(t) \end{pmatrix} = \bar{A}^p \hat{x}(t) + \bar{\Delta}_p Z^{t+p-1,k}. \tag{7}$$

By definition, $\bar{A} = A - KC$ is a stable matrix, so that for large p, $\bar{A}^p \approx 0$ and

$$\hat{x}(t+p) = \bar{\Delta}_p Z^{t+p-1,k}. \tag{8}$$

Substituting this result into the measurements equation,

$$y(t+p) = C\bar{\Delta}_p Z^{t+p-1,k} + Du(t+p) + e(t+p). \tag{9}$$

and propagating the output prediction for a window of length f, a set of equations can be stacked as

$$\begin{aligned} \begin{bmatrix} y(t+p) & \cdots & y(t+p+f) \end{bmatrix} &= C\bar{\Delta}_p \begin{bmatrix} Z^{t+p-1,k} & \cdots & Z^{t+p+f-1,t+f} \end{bmatrix} \\ &+ D \begin{bmatrix} u(t+p) & \cdots & u(t+p+f) \end{bmatrix} + \begin{bmatrix} e(t+p) & \cdots & e(t+p+f) \end{bmatrix}. \end{aligned} \tag{10}$$

Thus, estimates of $C\bar{\Delta}_p$ and D can be obtained as the solution of an unconstrained least squares problem. Let the controllability-like matrix $\bar{\Gamma}_f$ be defined as

$$\bar{\Gamma}_f = \begin{bmatrix} C \\ C\bar{A} \\ \vdots \\ C\bar{A}^{f-1} \end{bmatrix}. \tag{11}$$

The Hankel matrix $\bar{\Gamma}_f \bar{\Delta}_p$ is thus given by

$$\bar{\Gamma}_f \bar{\Delta}_p = \begin{bmatrix} C\bar{B} & C\bar{A}\bar{B} & \cdots & C\bar{A}^{p-1}\bar{B} \\ C\bar{A}\bar{B} & C\bar{A}^2\bar{B} & \cdots & C\bar{A}^{p}\bar{B} \\ \vdots & & \ddots & \\ C\bar{A}^{f-1}\bar{B} & C\bar{A}^{f}\bar{B} & \cdots & C\bar{A}^{p+f-1}\bar{B} \end{bmatrix}. \tag{12}$$

Given the approximation $\bar{A}^k \approx 0, k \geq p$,

$$\bar{\Gamma}_f \bar{\Delta}_p \approx \begin{bmatrix} C\bar{B} & C\bar{A}\bar{B} & \cdots & C\bar{A}^{p-1}\bar{B} \\ C\bar{A}\bar{B} & C\bar{A}^2\bar{B} & \cdots & 0 \\ \vdots & & \ddots & \\ C\bar{A}^{f-1}\bar{B} & C\bar{A}^{f}\bar{B} & \cdots & 0 \end{bmatrix}, \tag{13}$$

which can be constructed by blocks of the already computed matrix $C\bar{\Delta}_p$. Let $Z = \begin{bmatrix} Z^{t+p-1,k} & \cdots & Z^{t+p+f-1,t+f} \end{bmatrix}$. From the structure of Δ_p, it is possible to write

$$\hat{X} \doteq \begin{bmatrix} \hat{x}(t+p) & \cdots & \hat{x}(t+p+f) \end{bmatrix} = \bar{\Delta}_p Z. \tag{14}$$

Hence,

$$\bar{\Gamma}_f \hat{X} = \bar{\Gamma}_f \bar{\Delta}_p Z. \tag{15}$$

Since the right hand side is known, $\bar{\Gamma}_f \hat{X}$ can be computed. An infinite set of representations leads to the same system, as previously discussed. Thus, one might simply determine an $\hat{X}$ by performing a Singular Value Decomposition (SVD) such that

$$USV^T = \bar{\Gamma}_f \hat{X}. \tag{16}$$

Only a set of the largest $n_{\hat{x}}$ singular values might be considered, leading to a dimension $x \in \mathbb{R}^{n_{\hat{x}}}$. Thus, for instance, a compatible representation is given by

$$\bar{\Gamma}_f = U\sqrt{S}, \; \hat{X} = \sqrt{S}U^T. \tag{17}$$

From this point, the computation of A, B and C is straightforward. In particular, the system matrices can be determined by writing the innovation form dynamics in (3) for all inputs and measurements, in the same way as described in (10), using $\hat{X}$ as a known quantity. Note that K is an output of the PBSID.

The PBSID method has been proved to be a consistent estimator, even for closed-loop identification. Furthermore, as the implementation can be based on efficient and robust tools from linear algebra, it is extremely reliable numerically. Additionally, by choosing the number of singular values in the SVD, the size of the inner state x is a design variable. Finally, note that no initial guess is required for the identification of the system matrices A, B, C and D.

2.3 Virtual-Reference Predictor-Based tuning

The proposed VRPB approach builds upon the core idea of VRFT, that is, the construction of a virtual reference from data collected in a dedicated experiment, but replaces the direct regression of the control input on the virtual error with a prediction-based identification step based on PBSID. More precisely, in VRPB, the virtual reference and virtual error are defined exactly as in VRFT, ensuring that the design objective remains the solution of a model-reference control problem for the desired reference model $M(z)$. The solution however is instead constructed by means of a finite-horizon predictor 10 that captures the multi-step relationship between past virtual errors and control inputs, instead of fitting the control input u from the noisy virtual error e in a single least-squares step.

This enables the estimation of unstructured state-space controllers, allows the unified treatment of SISO and MIMO control problems and it provides a common framework for open- and closed-loop data. Finally, note that the SVD computed in (16) allows the user to choose the order of the controller in a data-driven sense.

The starting point is the same as in classical VRFT: from a data record

$$d_N = \{u(t), y(t)\}_{t=1,\dots,N}$$

and from the assigned reference model $M(z)$, a *virtual reference* $r_v(t)$ is defined so that

$$y(t) = M(z)r_v(t),$$

that is,

$$r_v(t) = M^{-1}(z)y(t),$$

whenever $M^{-1}(z)$ is well defined and stable, or after the standard filtering steps used in VRFT. The corresponding *virtual error* is then given by

$$e_v(t) = r_v(t) - y(t).$$

In the ideal case in which the measured data had been generated by the closed-loop system achieving exactly the desired model $M(z)$, the controller sought for would satisfy

$$u(t) = C(z)e_v(t).$$

Therefore, similarly to VRFT, the controller design problem can be recast as the problem of identifying, from data, a dynamical map from the virtual error e_v to the measured plant input u. The key difference with respect to standard VRFT is that this map is not estimated by a direct one-step regression of u on filtered versions of e_v, but rather through a predictor-based subspace identification step.

To this aim, the controller is described as an unknown linear dynamical system in innovation form,

$$\begin{aligned} x_c(t+1) &= A_c x_c(t) + B_c e_v(t) + K_c \varepsilon_c(t), \\ u(t) &= C_c x_c(t) + D_c e_v(t) + \varepsilon_c(t), \end{aligned} \tag{18}$$

where $x_c(t)$ is the controller state and $\varepsilon_c(t)$ is the innovation term associated with the identification problem. In other words, once the virtual signals have been constructed, the pair $(e_v(t), u(t))$ is interpreted as an input-output data set generated by the unknown controller dynamics.

At this point, the PBSID procedure described in Section 2.2 can be applied directly by considering $e_v(t)$ as the input and $u(t)$ as the output of the system to be identified. More precisely, one constructs an extended data vector

$$z_c(t) = \begin{pmatrix} e_v(t) \\ u(t) \end{pmatrix},$$

and uses a finite-horizon predictor to express future values of the controller state as a linear function of past virtual-error and input data. Following the same derivation leading to (10), one obtains a stacked predictor relation of the form

$$U_f = \Gamma_{c,f}\Delta_{c,p}Z_{c,p} + H_{c,f}E_f, \tag{19}$$

where U_f collects future values of u, $Z_{c,p}$ is built from past values of the extended signal z_c, $\Gamma_{c,f}$ is the observability-like matrix associated with the controller realization, and $\Delta_{c,p}$ contains the Markov parameters of the stable predictor. As in PBSID, least-squares and singular-value decomposition steps then provide an estimate of the controller state sequence and, subsequently, of the controller matrices (A_c, B_c, C_c, D_c).

This construction preserves the one-shot nature of VRFT, since the controller is still obtained directly from a single batch of data and from the assigned reference model. At the same time, it removes the need to postulate *a priori* a specific low-order controller structure such as PID or a fixed linearly parameterized transfer function. Indeed, VRPB identifies directly an unstructured state-space realization of the controller, whose order can be selected on the basis of the singular values arising in the SVD step. In this sense, the choice of controller complexity is performed in a data-driven way.

The use of PBSID within the virtual-reference framework also provides important practical advantages. First, the controller estimate is obtained from a dynamical multi-step predictor rather than from a static regression, which mitigates the adverse effect of the errors-in-variables nature of the classical VRFT regression. Second, the method naturally extends to MIMO problems, since the subspace machinery does not rely on scalar transfer-function parametrizations. Third, the predictor-based formulation is well suited to data collected under feedback, which is especially relevant in flight-control applications, where open-loop experiments are often impractical or unsafe.

Summarizing, the VRPB procedure can be described as follows:

1. collect an input-output data set $\{u(t), y(t)\}$ from the plant;
2. assign the desired closed-loop model $M(z)$;
3. compute the virtual reference $r_v(t)$ and the virtual error $e_v(t)$;
4. regard $(e_v(t), u(t))$ as the data generated by the unknown controller;
5. apply PBSID to identify a state-space realization of the controller;
6. select the controller order from the singular values and obtain the final controller realization.

It is important to remark that, as in VRFT, the quality of the result depends on the informativeness of the data set and on the compatibility between the assigned reference model and the achievable closed-loop behaviour. Moreover, the identified controller is optimal only in the implicit model-reference sense induced by the virtual-reference construction. Nevertheless, VRPB provides a flexible and practically appealing extension of VRFT, particularly when closed-loop data are used and when a structured controller parametrization is either unavailable or undesirable.

3 Simulation results

The aim of this section is to present preliminary results obtained by applying the VRPB approach to the tuning of an attitude control lopp for the longitudinal dynamics of a small-scale multirotor UAV.

3.1 Multirotor platform and control architectures

The considered multirotor platform, called ANT-X (see Figure 1), is a fixed-pitch quadrotor with the following characteristics:

- Take-Off Weight (TOW): approximately 250 grams;
- Flight time: 8 minutes;
- Frame dimensions (footprint): 200 mm (including rotors).

Fig. 1 The ANT-X UAV.

Concerning the baseline control architecture, the ANT-X platform adopts an attitude control scheme based on identical decoupled cascaded PID loops for the pitch, roll and yaw axes, running at 250 Hz. For instance, focusing on the pitch axis, the outer loop (measured angle ϑ, set-point ϑ^o) is a P controller, while the inner controller is a complete PID that computes the control torque M. More specifically, the derivative action of the inner loop is computed starting from the pitch rate q and not from the pitch angular rate error. On top of the pitch/pitch-rate cascade, an identical cascade structure deals with the longitudinal velocity/position control problems. Without loss of generality, in this work we will focus on the pitch axis alone.

As in most quadrotors, the decoupled architecture is justified by the fact that if the body axes are principal axes of inertia, then when the quadrotor is in near-hovering conditions the DOF could be assumed decoupled. Note, in passing, that the symmetric configuration of multirotors allows one to easily identify a set of principal axes of inertia. Hence, the attitude control problem is reduced to a set of three separate problems for, respectively, the pitch, roll and yaw axes. This, in turn, implies that the controllers for each axis can be tuned independently, for instance by exploiting data-driven such as the ones introduced in Section 2.1.

To this aim, note that the overall control structure can be converted into the regressor form in a straightforward step, resulting in:

$$u(t) = u(t-1) + \sum_{i=0}^{n_e} B_i e(t-i) + \sum_{j=0}^{n_y} B_j y(t-j) + \sum_{m=0}^{n_r} B_m r(t-m) \tag{20}$$

where the B_i are suitable scalars collecting the control parameters. While this structure leads to the VRFT-like approach, in the paper the more general unstructured VRPB approach is considered.

3.2 Simulation example: attitude control

In order to illustrate the application of the VRPB to a typical attitude control problem we consider the second-order model for pitch dynamics given by;

$$\dot{q} = M_q q + M_\delta \delta \tag{21}$$

$$\dot{\theta} = q \tag{22}$$

where the the input is the normalised pitching moment δ, we assume the measured output to be the pitch attitude θ and the derivatives M_q and M_δ are unknown parameters.

Simulation data is collected, in closed-loop, on the above model in order to generate the input and output datasets required for the VRPB algorithm. More precisely, a sine-sweep excitation sequence is applied. The input has a maximum excitation bandwidth of 50rad/s and an amplitude of 0.1. The latter is a non-dimensional amplitude, referred to the maximum moment that can be applied, consistently with the definition of the control input δ. The amplitude of the excitation signal has been selected in order to achieve a good signal-to-noise ratio without inducing nonlinear effect that would undermine the assumptions behind data-driven methods for LTI systems outlined in Section 2.

For the pitch loop, the reference model is a second order model, with continuous-time values for the desired bandwidth and damping ratio of 19rad/s and 0.7 respectively:

$$M(z) = \frac{0.002787z + 0.00269}{z^2 - 1.894z + 0.8991}.$$

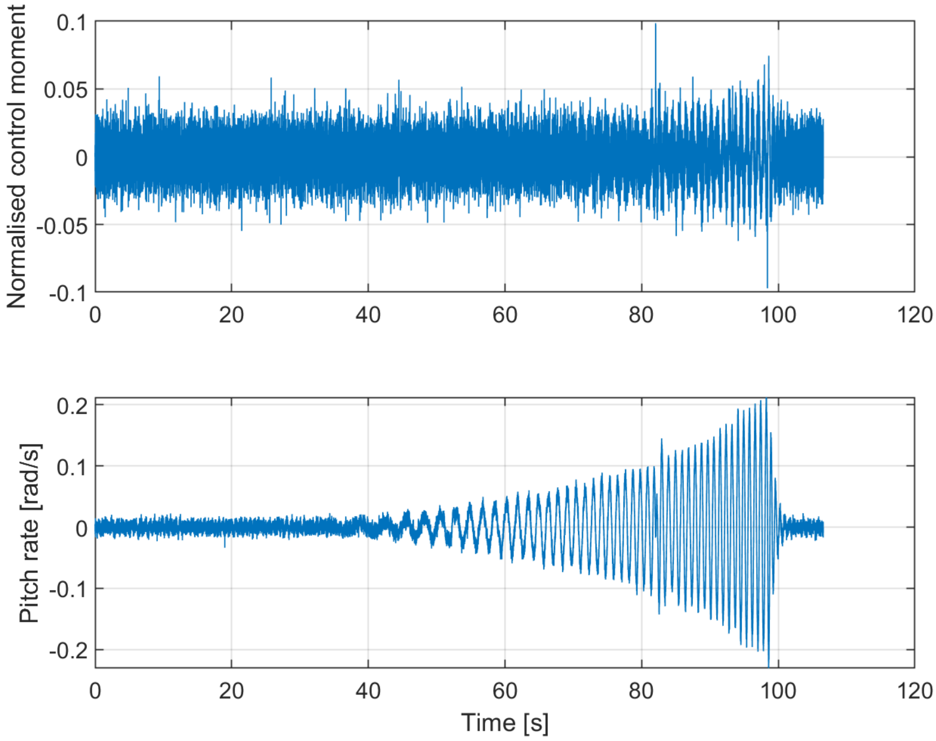

Fig. 2 Time histories of the input/output data used for VRPB tuning.

In this specific case no filtering action was needed, thus the weighting function has been defined as $W(z) = I$.

Following the application of the VRPB algorithm to the simulation data, with the hyperparameters p and f set to $p = f = 20$, the singular values depicted in Figure 3 have been obtained. Based on the inspection of the singular values, a second-order controller has been tuned. A comparison of the Bode plots and of the step responses obtained from the reference model and from the tuned closed-loop is presented in Figure 4 and in Figure 5. As can be seen from the figures, the match between the desired reference and the achieved closed-loop very satisfactory.

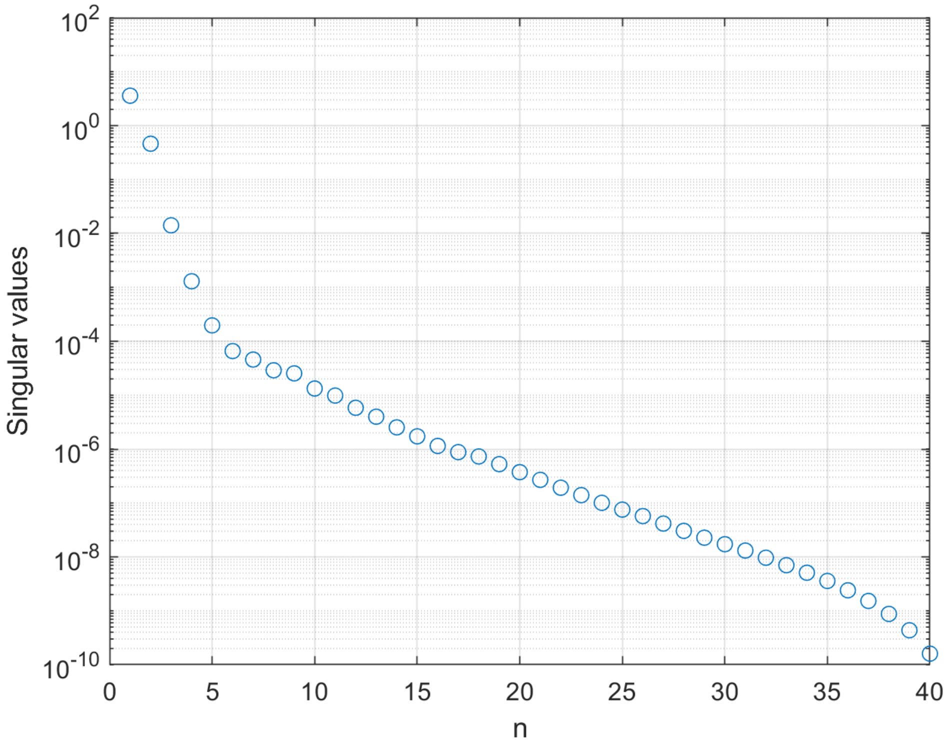

Fig. 3 Singular values of $\bar{\Gamma}_f \bar{\Delta}_p Z$

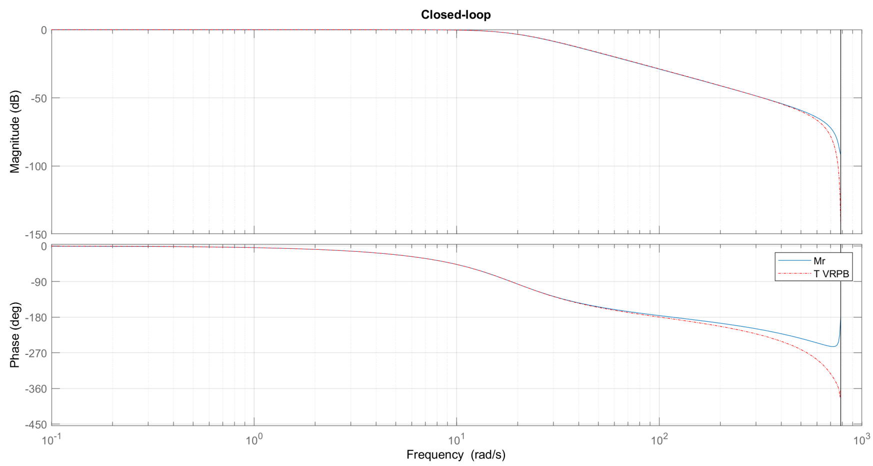

Fig. 4 Bode plots of the reference model and of the tuned complementary sensitivity

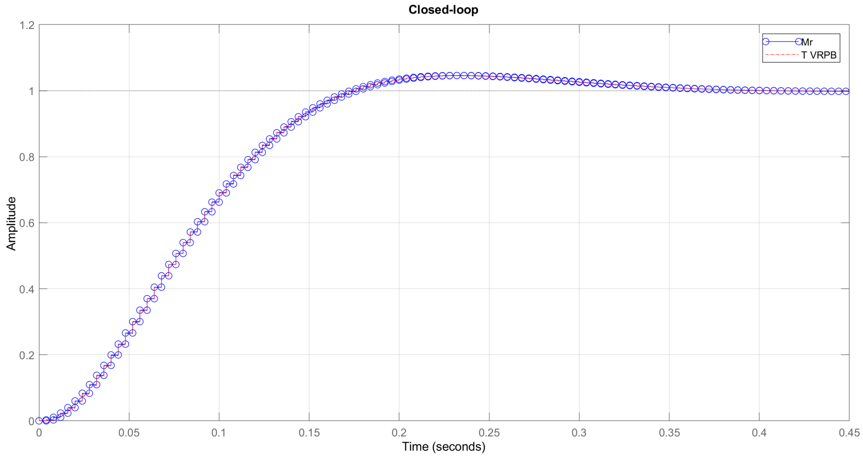

Fig. 5 Step responses of the reference model and of the tuned complementary sensitivity

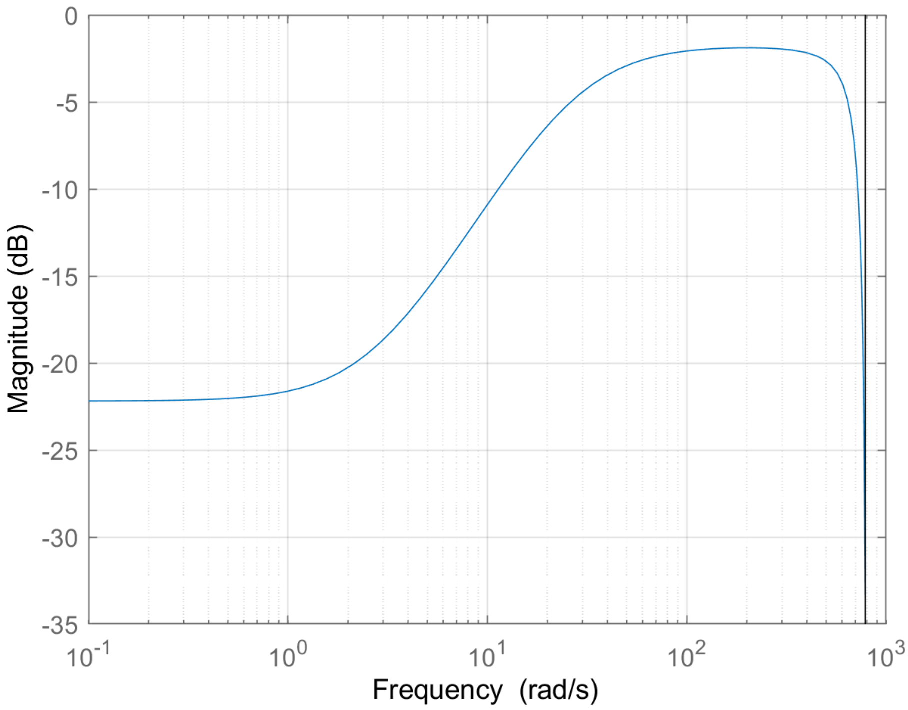

Fig. 6 Frequency-response function of the designed controller.

4 Conclusions

In this paper the VRPB method for data-driven controller tuning has beep proposed and its application to an attitude control problem for a multirotor UAV has been illustrated. The approach allows the non-iterative and systematic tuning of unstructured state-space controllers from experimental data, in a model-reference framework.

The application to a multirotor attitude control problem leads to satisfactor results in terms of matching between the desired reference model and the achieved closed-loop complementary sensitivity function.

Declaration of Use of Artificial Intelligence

AI was used in proofreading portions of the manuscript.

References

[1] R. Mahony, V. Kumar, and P. Corke. Multirotor Aerial Vehicles: Modeling, Estimation and Control of Quadrotor. *IEEE Robotics & Automation Magazine*, 19(3):20–32, 2012.

[2] T. P. Nascimento and M. Saska. Position and attitude control of multi-rotor aerial vehicles: A survey. *Annual Reviews in Control*, 48:129–146, 2019.

[3] F. Riccardi, P. Panizza, and M. Lovera. Identification of the attitude dynamics for a variable-pitch quadrotor UAV. In *40^{th} European Rotorcraft Forum, Southampton, UK*, pages 1–9, 2014.

[4] H. Hjalmarsson, M. Gevers, S. Gunnarsson, and O. Lequin. Iterative feedback tuning: theory and applications. *IEEE Control Systems*, 18(4):26–41, 1998.

[5] 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.

[6] K. Van Heusden, A. Karimi, and D. Bonvin. Data-driven model reference control with asymptotically guaranteed stability. *International Journal of Adaptive Control and Signal Processing*, 25(4):331–351, 2011.

[7] S. Formentin, K. 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.

[8] L. Campestrini, M. Gevers, and A.S. Bazanella. Virtual Reference Feedback Tuning for Non Minimum Phase Plants. In *European Control Conference (ECC 2009), Budapest, Hungary, pp. 1955-1960*, 2009.

[9] S. Formentin, M. Corno, S.M. Savaresi, and L. Del Re. Direct data-driven control of linear time-delay systems. *Asian Journal of Control*, 13(5):1–12, 2011.

[10] S. Formentin and A. Karimi. Enhancing statistical performance of data-driven controller tuning via L_2-regularization. *Automatica*, 50(5):1514–1520, 2014.

[11] F. Previdi, T. Schauer, S.M. Savaresi, and K.J. Hunt. Data-driven control design for neuroprotheses: a virtual reference feedback tuning (VRFT) approach. *IEEE Transactions on Control Systems Technology*, 12(1):176–182, 2004.

[12] S. Formentin, M.C. Campi, and S.M. Savaresi. Virtual reference feedback tuning for industrial PID controllers. In *19th IFAC World Congress, Cape Town, South Africa*, pages 11275–11280, 2014.

[13] A. Zangarini, D. Invernizzi, P. Panizza, and M. Lovera. Closed-loop MIMO data-driven attitude control design for a multirotor UAV. *CEAS Aeronautical Journal*, 11(4):873–884, 2020.

[14] A. Chiuso and G. Picci. Consistency analysis of some closed-loop subspace identification methods. *Automatica*, 41(3):377–391, 2005.

[15] van der Veen, G., van Wingerden, J.-W., Bergamasco, M., Lovera, M., Verhaegen, M. Closed-loop subspace identification methods: an overview. *IET Control Theory & Applications*, 7(10):1339–1358, 2013.

[16] I. Houtzager. *Subspace Identification in Closed-loop and Applications to Wind Turbine Control*. PhD thesis, Delft University of Technology, 2010.

[17] I. Houtzager, J. W. van Wingerden, and M. Verhaegen. VARMAX-based closed-loop subspace model identification. *Automatica*, 49(11):3122–3130, 2013.

[18] A. Chiuso. The role of vector autoregressive modeling in predictor-based subspace identification. *Automatica*, 43(6):1034–1048, 2007.

[19] S. Formentin, A. Cologni, D. Belloli, F. Previdi, and S.M. Savaresi. Fast tuning of cascade control systems. In *18th IFAC World Congress, Milan, Italy*, pages 10243–10248, 2011.

[20] D. Invernizzi, P. Panizza, F. Riccardi, S. Formentin, and M. Lovera. Data-driven attitude control law of a variable-pitch quadrotor: a comparison study. *IFAC-PapersOnLine*, 49(17):236 – 241, 2016. ISSN: 2405-8963. 20th IFAC Symposium on Automatic Control in Aerospace 2016.