

A Learn-to-Fly Approach for Adaptively Tuning Flight Control Systems

Jared A. Grauer

Received: date / Accepted: date

Abstract A method is presented for adaptively tuning feedback control gains in a flight control system to achieve desired closed-loop performance. The method combines efficient aircraft parameter estimation, for identifying closed-loop dynamics models, with online nonlinear optimization, for sequentially perturbing and updating control gains to improve performance. Access to flight measurements and the control gains is required, but no prior information about the aircraft dynamics or the flight control architecture is needed. After the procedure, the optimized control gains (with uncertainties), the open-loop dynamics model, and the closed-loop dynamics model are available. The method is demonstrated for tuning a longitudinal stability augmentation system using a realistic nonlinear flight dynamics simulation of a subscale airplane. Convergence was attained using five maneuvers that spanned approximately one

Jared A. Grauer

NASA Langley Research Center

Hampton, VA

minute of flight test time. Although shown for a relatively simple case, the method is general and can be applied to other aircraft, axes of motion, performance metrics, and control system designs.

Keywords Adaptive tuning · Real-time system identification · Flight control

Mathematics Subject Classification (2000) 49J53 · 49K99 · more

1 Introduction

Conventional methods for designing control laws and choosing feedback gains in an aircraft flight control system can include first-principles and computational analyses, wind tunnel testing, simulation development, Monte Carlo investigations, and flight test verification. Tuning the control gains can be an iterative process involving several of these analyses, each of which may require significant time and cost. It can therefore be a considerable undertaking to tune flight control systems for adequate closed-loop behavior.

In an effort to shift the current paradigm for flight control design, the NASA Learn-to-Fly project [1,2] has been developing algorithms for online system identification, guidance, and control that adapt in real time to measured data without using prior knowledge from conventional ground-based analyses. One practical application of this concept is to provide tools to control engineers for rapidly tuning aircraft control systems during flight using models identified in real time. This concept bypasses iterations of the aforementioned conventional process.

This paper presents a feasibility study for one such solution. The approach is based on an online optimization for the control gains that admit the desired closed-loop performance, as determined by models identified using system identification. The approach is based on flight test data and does not use prior information about the aircraft dynamics or workings of the flight control architecture. The approach is general and can be applied to any set of dynamics and to any flight control system having adjustable control gains.

The organization of this paper is as follows. Section 2 describes the proposed method to adaptively tune the control gains. Section 3 presents an example where control gains in a longitudinal stability augmentation system are tuned to achieve a desired damping and natural frequency for a realistic nonlinear airplane simulation. Computer software used in the analysis is available in a MATLAB[®] toolbox called System IDentification Programs for AirCraft [3], which is associated with Ref. [4]. Section 4 summarizes the findings and concludes the discussion.

2 Adaptive Tuning Method

The goal of the adaptive tuning method is to experimentally determine the control gains in a flight control system that meet some specified metrics for closed-loop performance. The approach is to conduct a nonlinear optimization during flight to optimize the control gains, based on the measured performance of identified closed-loop models. Specifically, the process begins by exciting the aircraft about a reference flight condition to identify a baseline




Fig. 1 Algorithm flow chart for adaptive tuning of flight control gains.

closed-loop model and measure the closed-loop performance. Afterwards, each control gain is perturbed to detect changes in the closed-loop performance. A nonlinear optimization step is then used to update the control gains based on this information, and the process is repeated until results converge to within acceptable levels. A flow chart depicting this process is illustrated in Fig. 1.

In terms of Ref. [5], the approach can be considered a direct method of adaptive control because control gains are directly optimized for closed-loop performance, although dynamic models are also identified. The approach is similar to extremum control or peak-seeking control [5,6], which has been

used to reduce drag of an aircraft in formation flight [7,8] and to optimize trim settings of an F-18 for minimum fuel consumption [9,10]. The novel parts of the proposed method are the optimization implementation, the inclusion of efficient aircraft system identification methods, and the applications.

Details of the proposed technique are discussed in the following sections. First, the structure of the postulated dynamic model describing the closed-loop system is presented. Second, the method for estimating parameters within this model from the measured flight data is summarized. Last, the nonlinear optimization for the control gains is described in detail.

2.1 Postulated Dynamic Model

A model of the aircraft flight dynamics is needed to perform system identification and extract a metric of closed-loop performance. From the point of view of the pilot, most aircraft operating in typical flight conditions and over a certain frequency bandwidth of interest behave similarly to classical models of flight dynamics [11]. In these cases, a Low-Order Equivalent System (LOES) model [12,13] is appropriate for identification of the closed-loop dynamics. Otherwise, other closed-loop models could be substituted.

The decoupled longitudinal short-period approximation used later in Section 3 has the LOES state-space perturbation model

$$\begin{bmatrix} \dot{\alpha}(t) \\ \dot{q}(t) \end{bmatrix} = \begin{bmatrix} -L_\alpha & 1 - L_q \\ M_\alpha & M_q \end{bmatrix} \begin{bmatrix} \alpha(t) \\ q(t) \end{bmatrix} + \begin{bmatrix} -L_\eta \\ M_\eta \end{bmatrix} \Delta\eta(t - \tau) \quad (1)$$

The input is the pilot longitudinal stick displacement η , and the states are angle of attack α and pitch rate q . The equivalent time delay τ approximates latencies due to actuator lags, the control system, nonlinearities, higher-order dynamics, and other factors. Both the inputs and states represent small deviations from the reference flight condition. Terms inside the state matrices are the stability and control derivatives (e.g., L_α) and are assumed to be constants for a given flight condition.

Typically, Eq. (1) is used to describe the bare-airframe dynamics of the aircraft, where the inputs are the control surfaces and the stability and control derivatives describe the aerodynamic forces and moments. In this application, the inputs are the pilot commands, and the stability and control derivatives subsume the closed-loop system. Section 2.2 describes the identification of the unknown closed-loop stability and control derivatives.

The block diagram in Fig. 2 shows the generalized subsystems enclosed by the LOES model, which is shaded in gray. The LOES model contains the actuators, airplane bare-airframe dynamics, sensors, and flight control system. The control block contains the flight control system, which has the control gains to be tuned to meet closed-loop performance objectives. The effect of these control gains are subsumed into the identified stability and control derivatives in the LOES model.




Fig. 2 Block diagram of the closed-loop LOES model.

2.2 Parameter Estimation

Parameter estimation was used to determine values and uncertainties for unknown parameters in Eq. (1). Because correlation problems preclude the simultaneous identification of both the time delay and the stability and control derivatives [13], a relaxation technique was employed to first estimate the time delay and then hold it fixed at that estimate for the remaining analysis.

The time delay was determined using measurements of the pilot stick and the elevator surface deflection. In the frequency domain, the model relating these quantities was [14]

$$\tilde{\delta}(j\omega) = \tilde{\eta}(j\omega) e^{-j\omega\tau} \quad (2)$$

where $\tilde{\eta}(j\omega)$ denotes the finite Fourier transform of $\eta(t)$ for frequencies ω over a bandwidth of interest. This is a nonlinear estimation problem that can be solved using output error [4], which is closely related to the control gain optimization discussed next in Section 2.3. Starting the optimization with an

initial estimate $\hat{\tau} = 0$ is typically sufficient and converges quickly because the optimization is problem well conditioned. This estimation was performed once at the beginning of the tuning procedure and was then held constant for the remaining analysis because the equivalent time delay characterizes latencies that are not expected to significantly change between maneuvers.

With the time delay determined and fixed, the estimation of the closed-loop stability and control derivatives can decouple into two linear estimation problems, which can be efficiently solved using equation error with frequency-domain data [4]. The model equations are

$$-\frac{g}{V_0} \tilde{a}_z(j\omega) = L_\alpha \tilde{\alpha}(j\omega) + L_q \tilde{q}(j\omega) + L_\eta \tilde{\eta}(j\omega) e^{-j\omega\hat{\tau}} \quad (3a)$$

$$j\omega \tilde{q}(j\omega) = M_\alpha \tilde{\alpha}(j\omega) + M_q \tilde{q}(j\omega) + M_\eta \tilde{\eta}(j\omega) e^{-j\omega\hat{\tau}} \quad (3b)$$

where a_z are vertical accelerometer measurements at the aircraft mass center, g is the acceleration due to gravity, and V_0 is the trim airspeed. These equations are appropriate for small-perturbation maneuvers about a reference flight condition at low angles of attack.

Each part of Eq. (3) is of the form

$$\tilde{\mathbf{z}} = \tilde{\mathbf{X}}\boldsymbol{\theta} + \tilde{\mathbf{v}} \quad (4)$$

where $\mathbf{z}$ is the modeled variable, $\mathbf{X}$ is the regressor matrix of explanatory variables, $\boldsymbol{\theta}$ is a vector of unknown model parameters, and $\mathbf{v}$ is the equation error. The ordinary least-squares solution to Eq. (4) is

$$\hat{\boldsymbol{\theta}} = \left[\Re \left(\tilde{\mathbf{X}}^\dagger \tilde{\mathbf{X}} \right) \right]^{-1} \Re \left(\tilde{\mathbf{X}}^\dagger \tilde{\mathbf{z}} \right) \quad (5)$$

where $\Re$ returns the real part of the argument and where † is the complex conjugate transpose. The uncertainty in the parameter estimates is reflected by the covariance matrix

$$\text{cov}(\hat{\boldsymbol{\theta}}) = \hat{\sigma}^2 \left[\Re \left(\tilde{\mathbf{X}}^\dagger \tilde{\mathbf{X}} \right) \right]^{-1} \quad (6)$$

where the error variance $\hat{\sigma}^2$ is computed from the model residuals [4]. These equations are processed twice: once with Eq. (3a) for the lift force derivatives, and once with Eq. (3b) for the pitching moment derivatives.

Estimating the stability and control derivatives in this manner using equation error produces a solution without iteration, which is important for in-flight control gain tuning. Because the estimation is performed in the frequency domain over a bandwidth of interest where the signal-to-noise ratios are high, and for other reasons, the identified parameters and the associated uncertainties are statistically accurate and exhibit some robustness to measurement noise and disturbances [4].

2.3 Control Gain Optimization

This section describes how the control gains are optimized during flight based on the closed-loop performance metrics extracted from identified closed-loop models. The desired set of performance metrics are denoted by the vector $\boldsymbol{\gamma}_d$. The error in the closed-loop performance is

$$\boldsymbol{\nu} = \boldsymbol{\gamma}_d - \boldsymbol{\gamma} \quad (7)$$

where $\boldsymbol{\gamma}$ are the performance metrics computed from the identified closed-loop models. The goal of the optimization is to determine the vector of control gains $\mathbf{k}$ that minimizes the negative-log maximum likelihood cost function

$$J(\mathbf{k}) = \frac{1}{2} \boldsymbol{\nu}^T \mathbf{R}^{-1} \boldsymbol{\nu} + \frac{1}{2} \ln |\mathbf{R}| \quad (8)$$

In Eq. (8), the weighting matrix $\mathbf{R}$ is the error covariance matrix, which quantifies uncertainty on the computed performance metrics. To obtain $\mathbf{R}$, the unscented transform [15] was used to propagate $\text{cov}(\hat{\boldsymbol{\theta}})$ from Eq. (6) to $\text{cov}(\boldsymbol{\nu})$, which involves a nonlinear transform. Comparison against Monte Carlo simulations for the example in Section 3 confirmed the unscented transformed produced accurate error covariance estimates in a computationally efficient manner for this purpose.

Equation (8) admits a nonlinear optimization problem that requires iteration from starting values. The optimization update for the i^{th} iteration was computed using the Gauss-Newton method as

$$\hat{\mathbf{k}}_i = \hat{\mathbf{k}}_{i-1} - \mathbf{M}^{-1} \frac{\partial J}{\partial \mathbf{k}} \quad (9)$$

where

$$\mathbf{M} = \frac{\partial \boldsymbol{\gamma}^T}{\partial \mathbf{k}} \mathbf{R}^{-1} \frac{\partial \boldsymbol{\gamma}}{\partial \mathbf{k}} \quad (10)$$

is the information matrix and where

$$\frac{\partial J}{\partial \mathbf{k}} = -\frac{\partial \boldsymbol{\gamma}^T}{\partial \mathbf{k}} \mathbf{R}^{-1} \boldsymbol{\nu} \quad (11)$$

is the local cost gradient. The uncertainties in the control gain estimates are

$$\text{cov}(\hat{\mathbf{k}}) = \mathbf{M}^{-1} \quad (12)$$

In Eqs. (9)–(11), the performance sensitivity matrix $\frac{\partial \gamma}{\partial \mathbf{k}}$ quantifies how much each control gain changes each performance metric. This matrix was determined by finite differences, as depicted by the inner loop in Fig. 1, wherein each control gain was perturbed, a maneuver performed, and a new closed-loop model identified. The control gain perturbations should be small enough that the local cost gradient is accurately approximated, but must be large enough that statistical differences between identified closed-loop models can be discerned. For the example in Section 3, perturbations were selected as 10% of the current control gain value, and backward one-sided finite differences were used.

The optimization procedure was considered converged when all of the performance metrics were achieved within a threshold amount of the desired values. For the example in Section 3, a 2% threshold was used.

Nonlinear optimizations can exhibit slow convergence or divergence if the starting values for the optimization are not sufficiently close to the unknown optimal solution. To start the optimizer near the optimized solution, pole placement techniques or control law designs based on prior models could be used for the first iteration of the control gain optimization.

3 Example

3.1 Nonlinear Aircraft Simulation

In this example, the nonlinear flight dynamics simulation of a Hanger 9 Ultra-StickTM 120, developed under the NASA Free-flying Aircraft for Sub-scale Experimental Research (FASER) project, was used. The augmented research airplane for which the simulation is based has a 6.3 ft wingspan and 16.4 lbf weight, and is shown in Fig. 3. The simulation includes the nonlinear six-degree-of-freedom equations of motion and an aerodynamic model that uses multivariate polynomial splines generated from wind tunnel data. Modeled inputs for the simulation are the conventional throttle, elevator, aileron, and rudder surface deflections. See Refs. [16,17] for complete details on the simulation.




Fig. 3 The FASER test aircraft (credit: Ref. [18]).

The nonlinear simulation was modified in three ways. First, the longitudinal stability augmentation system control law

$$\delta_c = \eta - k_\alpha \alpha - k_q q \quad (13)$$

was included, which adds proportional feedback from measurements of α and q to the pilot commands. Second, actuator dynamics were included as first-order filters with time delays. Lastly, realistic noise sequences for the aircraft were applied to the simulated measurements.

3.2 Results

The nonlinear simulation was trimmed for steady, wings-level flight at 80 ft/s airspeed, 1000 ft altitude, and 2.5 deg angle of attack. The open-loop short period mode at this condition was stable and had damping ratio 0.8 and natural frequency 7.3 rad/s (1.2 Hz).

The goal for tuning was to determine the feedback gains $\mathbf{k} = [k_\alpha, k_q]^T$ in the stability augmentation control law given by Eq. (13) to change the short period damping and frequency to 0.7 and 5.0 rad/s (0.8 Hz), respectively. Therefore, the metric of closed-loop performance was $\gamma_d = [\zeta, \omega_n]^T = [0.7, 5.0]^T$. The purpose was to slow the response and make the aircraft easier to fly by a remote pilot, who might be contending with visibility issues and other factors during flight testing.

The presented example used five discrete excitation maneuvers (65 s of total modeling data) to determine the feedback gains that achieved the desired closed-loop short period damping and frequency within 2%. Figure 4 shows the concatenated modeling data. For each maneuver, an identical small-amplitude input with frequency content near the short period mode was applied by the pilot for 10 s. After each maneuver, the tuning method changed the control

gains and the corresponding elevator and aircraft responses changed in shape and amplitude.




Fig. 4 Measurements for five concatenated excitation maneuvers.

The migration of eigenvalues for the identified closed-loop models is shown in Fig. 5. Poles 1 were for the first maneuver, in which the airplane was flown open loop. These data were also used to determine $\hat{\tau}$, which was held constant for the remaining analysis. Based on the first identified model, pole placement was initially used to update the control gains. Poles 2 were closer to the desired location, but remained outside the 2% threshold due to nonlinearities, parametric and unstructured errors in the identified model, and random noise in the measured data. Poles 3 and 4 were the result of sequentially reducing

k_α and k_q , respectively, by 10%. Poles 5, which were inside the desired 2% boundary, were attained after computing the local cost gradient at Poles 2 using Poles 3 and 4, and taking one Gauss-Newton update step.




Fig. 5 Closed-loop pole migration during tuning (positive imaginary axis only).

Parameter estimates using measured data from the five maneuvers are shown in Fig. 6. In general, the estimates were reasonable, had relatively low uncertainties (which are shown as $\pm 2\sigma$ bounds), and were in statistical agreement with numerical linearizations of the closed-loop simulation. Based on the first maneuver of the open-loop dynamics, the equivalent time delay was identified as $\hat{\tau} = 0.049 \pm 0.001$ s. Expected changes to the estimated M_α and M_q closed-loop derivatives resulting from perturbations to the k_α and k_q control gains were exhibited in the estimates for maneuvers 3 and 4, respectively.

Figure 7 shows the evolution of the performance metrics over the five maneuvers, and Fig. 8 shows the resulting control gains. The converged solution was $\hat{\mathbf{k}} = [-0.038 \pm 0.046, 0.044 \pm 0.005]^T$, which resulted in closed-loop performance $\hat{\gamma} = [0.690 \pm 0.038, 4.990 \pm 0.180]^T$. The smaller values and uncertainties




Fig. 6 Identified closed-loop model parameters and $\pm 2\sigma$ error bounds.

for k_q were a result of the system having higher sensitivities to that gain. For the open-loop system identified using data from the first maneuver, the performance sensitivity matrix was

$$\frac{\partial \gamma}{\partial \mathbf{k}} = \begin{bmatrix} \frac{\partial \zeta}{\partial k_\alpha} & \frac{\partial \zeta}{\partial k_q} \\ \frac{\partial \omega_n}{\partial k_\alpha} & \frac{\partial \omega_n}{\partial k_q} \end{bmatrix} = \begin{bmatrix} +0.40 & -0.85 \\ -4.00 & -36.0 \end{bmatrix} \quad (14)$$

which indicated that k_q was more effective at changing the closed-loop damping and frequency than k_α , which was statistically zero.

3.3 Further Discussion

The simulation results shown in Section 3.2 were a baseline result. Different parts of the tuning process were changed using this example to investigate robustness of the method.

One aspect of the analysis investigated was the finite differences used to compute $\frac{\partial \gamma}{\partial \mathbf{k}}$. Backward differences, where the control gains perturbations decreased the gains, were preferred to forward differences as a safe-guard against




Fig. 7 Estimated closed-loop performance metrics and $\pm 2\sigma$ error bounds.




Fig. 8 Estimated control gains and $\pm 2\sigma$ error bounds.

large aircraft responses. Central finite differences resulted in more accurate sensitivity matrices but on average used more maneuvers and flight time to converge.

The control tuning process could have also begun from non-zero control gains, instead of starting the aircraft in open loop. This would be the case where ground-based predictions of control gains were available, or for when

perturbations to the nominal control laws were requested for other flight conditions, or to account for aircraft aging as a maintenance function.

The parameter estimation method in Section 2.2 assumed a linear time-invariant system, which necessitated changing the control gains in a discrete manner. If a model for identification could be postulated that explicitly captures the nonlinear effect of the control gains on the performance metrics, then the control gains could be changed in a continuous manner. Orthogonal functions could be used to change all the gains simultaneously, leading to short durations of flight data needed for convergence of the control gains.

Other performance metrics could have been chosen for the optimization than the ones used in Section 3. For example, the same closed-loop poles could have been specified as real and imaginary parts, or rather as settling time and overshoot. The only changes would be to implement the way in which these performance metrics are obtained from the identified closed-loop models. Also, other convergence criteria could have been used, depending on the application requirements. For example, when the 95% uncertainty fell within the 2% threshold and the desired performance was obtained in a statistical sense, the optimization could have been stopped. In the presented example, this would have occurred after the fourth maneuver was processed, according to Fig. 7.

Although this example was presented for a relatively simple control law and dynamic system, the approach is general and can be extended to other axes, other models, and more complicated control laws. It is only required

that a closed-loop model can be identified and that the control gains can be adjusted during flight. However, as the number control gains increases, the number of maneuvers and flight time duration necessary to tune those gains also increases.

4 Conclusions

A direct method was presented for adaptively tuning control gains in a flight control system to achieve specified closed-loop performance metrics within defined tolerances. The method consists of sequentially perturbing control gains about nominal values and computing performance metrics from identified closed-loop models. Cost gradients were computed and an online nonlinear optimization was used to update the nominal control gain values using a Gauss-Newton step to improve performance. This process was repeated during flight until the desired closed-loop performance was achieved. The method was demonstrated using a nonlinear flight dynamics simulation of a subscale airplane to tune a longitudinal stability augmentation system and achieve desired closed-loop poles.

It was found using the nonlinear simulation that the approach is feasible. Relatively short amounts of time and numbers of maneuvers were needed to converge on control gains to acquire desired damping and frequency of the closed-loop short period mode. No previous modeling data or knowledge of the internal workings of the control system were used. The method provided

uncertainties on the estimated control gains, which provide insight into the efficacy of the designed control laws.

The requirements to perform this automated tuning process are determining a model structure that adequately reflects the closed-loop dynamics and can be used for identification, and having access to the tunable gains during flight. LOES models were used in the simulation results presented, but there may be other situations where these models are not appropriate. Although the tuning method can be performed with any aircraft axis and any arbitrary control law, the number of maneuvers needed and the corresponding flight test time increases with the number of control gains. For complex control laws, this could take significant amounts of flight time, potentially taking the airplane off the desired flight condition.

The main applications conceived were for small aircraft manufacturers or homebuilders, and for avionics designers who wish to design and tune control laws without the ground-based design and iteration that is usually incurred. This method could also be used as a way to rapidly tune controllers during envelope-expansion flight of research aircraft, based on computer simulation models. The method could also be used by smaller research outfits with sub-scale aircraft, trying to obtain stabilized flight platforms with which to conduct subsequent flight research.

Acknowledgements This work was funded by the NASA Aeronautics Research Mission Directorate (ARMD) through the Transformational Tools and Technologies (TTT) project. Discussions with Eugene Morelli, Jack Ryan, and members of the NASA Learn-to-Fly project

are acknowledged and appreciated. Lastly, the author would like to acknowledge the encouragement of Eugene Heim in performing this work.

References

1. Heim, E., Viken, E., Brandon, J., Croom, M.: NASA's Learn-to-Fly Project Overview. In: Atmospheric Flight Mechanics Conference, AIAA Paper 2018-3307 (Jun. 2018). DOI 10.2514/6.2018-3307
2. Riddick, S.: An Overview of NASA's Learn-to-Fly Technology Development. In: Atmospheric Flight Mechanics Conference, AIAA Paper 2020-0760 (Jan. 2020). DOI 10.2514/6.2020-0760
3. Morelli, E.: System IDentification Programs for AirCRAFT (SIDPAC). <https://software.nasa.gov/software/LAR-16100-1>. Accessed: Jan. 22, 2020
4. Morelli, E., Klein, V.: Aircraft System Identification: Theory and Practice, second edn. Sunflyte, Williamsburg, VA (2016)
5. Åström, K., Wittenmark, B.: Adaptive Control, second edn. Addison Wesley, Reading, MA (1995)
6. Ariyur, K., Krstić, M.: Real-Time Optimization by Extremum-Seeking Control. Wiley, Hoboken, NJ (2003)
7. Ryan, J., Speyer, J.: Peak-Seeking Control Using Gradient and Hessian Estimates. In: American Control Conference, WeA17.5 (Jun. 2010). DOI 10.1109/ACC.2010.5531081
8. Hanson, C., Ryan, J., Allen, M., Jacobson, S.: An Overview of Flight Test Results for a Formation Flight Autopilot. In: Guidance, Navigation, and Control Conference, AIAA Paper 2002-4755 (Jun. 2002). DOI 10.2514/6.2002-4755
9. Schaefer, J., Brown, N.: Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Architecture and Performance Metrics. In: Guidance, Navigation, and Control Conference, AIAA Paper 2013-4856 (Aug. 2013). DOI 10.2514/6.2013-4856
10. Brown, N., Schaefer, J.: Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Flight-Test Results. In: Guidance, Navigation, and Control Conference, AIAA Paper 2013-5171 (Aug. 2013). DOI 10.2514/6.2013-5171
11. McRuer, D., Ashkenas, I., Graham, D.: Aircraft Dynamics and Automatic Control. Princeton University Press, Princeton, NJ (1973)

12. Anon.: Flying Qualities of Piloted Aircraft. Tech. Rep. Mil-Std 1797a, Department of Defense (1990)
13. Morelli, E.: Identification of Low Order Equivalent System Models from Flight Test Data. NASA TM-2000-210117 (Aug. 2000)
14. Morelli, E.: Dynamic Modeling from Flight Data with Unknown Time Skews. *Journal of Guidance, Control, and Dynamics* **40**(8), 2083–2091 (2017)
15. Uhlmann, J.: Dynamic Map Building and Localization: New Theoretical Foundations. Ph.D. thesis, University of Oxford (1995)
16. Morelli, E.: A Collection of Nonlinear Aircraft Simulations in MATLAB. NASA TM-2003-212145 (Jan. 2003)
17. Morelli, E.: A Collection of Nonlinear Aircraft Simulations in MATLAB. <https://software.nasa.gov/software/LAR-17463-1>. Accessed: Jan. 22, 2020
18. Denham, C., Owens, D.: Rotary Balance Wind Tunnel Testing for the FASER Flight Research Aircraft. In: Atmospheric Flight Mechanics Conference, AIAA Paper 2016-3105 (Jun. 201). DOI 10.2514/6.2016-3105