

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

Jared A. Grauer*

NASA Langley Research Center, Hampton, Virginia 23681

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 parameter estimation for identifying closed-loop dynamics models, with online nonlinear optimization for sequentially perturbing and updating control gains to improve performance. Prior information on stability and control derivatives is not needed, nor is any knowledge about the control system architecture. Following convergence, 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 the NASA FASER airplane. Convergence was attained using five piloted maneuvers that spanned approximately one minute of flight test time. Although demonstrated for a relatively simple case, the method is general and can be applied to other aircraft, axes, performance metrics, and control systems.

Nomenclature

b	wing span, ft	Δ	perturbation
\bar{c}	mean aerodynamic chord, ft	δ_t	throttle setting
∂	partial derivative	ω	frequency, rad/s
$\text{cov}(\cdot)$	covariance	$ \cdot $	absolute value
g	gravitational acceleration, ft/s ²		
h	altitude, ft	Subscripts	
\Im	imaginary part	0	reference value
$I_{..}$	inertia components, slug-ft ²	Superscripts	
j	imaginary number, $\sqrt{-1}$	⁻¹	inverse
m	mass, slug	T	transpose
\Re	real part	\dagger	complex conjugate transpose
S	wing reference area, ft ²	\cdot	time derivative
s	Laplace variable	$\hat{\cdot}$	estimate
t	time, s		
x_{cm}	longitudinal mass center position, in		

*Research Engineer, Dynamic Systems and Control Branch, MS 308. Senior Member AIAA.

I. Introduction

THE NASA Learn-to-Fly project has proposed replacing the conventional flight control development process with automated and efficient flight-based algorithms, as depicted in Fig. 1. References [1–10] present advances in aircraft system identification and feedback control developed during the project towards this goal. The Learn-to-Fly concept can also be applied to the related problem of selecting proper feedback control gains in a flight control system to achieve desired closed-loop performance. An adaptive, real-time tuning method would allow airplane or avionics manufacturers to rapidly modify and verify flight control systems without the costs and durations incurred in the current paradigm. Airplane performance could also be improved by periodically adjusting the gains, as a maintenance function, to compensate for repairs or modifications, aircraft aging, or manufacturing differences between aircraft.

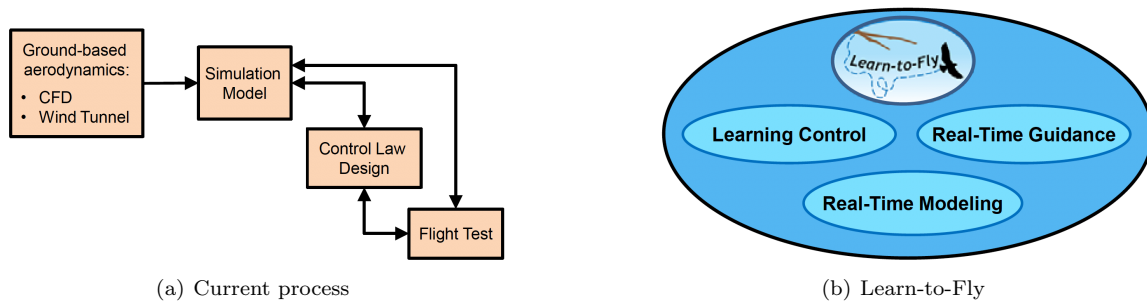


Figure 1. Approaches for flight control development, adapted from Ref. [1]

This paper presents a feasibility study for one such solution. The approach combines efficient system identification with online nonlinear optimization to iteratively perturb and update control gains until the desired performance is achieved. All information is derived from flight test data with no requirement for prior dynamic modeling or knowledge of the flight control system architecture. Standard flight test instrumentation is assumed, and excitation inputs for system identification can be provided by the pilot. The method is also robust to moderate levels of measurement noise, turbulence, and modeling error.

This paper is organized as follows. Section II describes the adaptive tuning method. In Section III, the nonlinear flight dynamics simulation of the NASA Free-flying Aircraft for Sub-scale Experimental Research (FASER) aircraft [11] is introduced. Application of the tuning method to a longitudinal stability augmentation system (SAS) for the FASER nonlinear simulation is demonstrated in Section IV. Lastly, conclusions are discussed in Section V. Although a relatively simple example is used in this paper, the method can be applied to other aircraft, axes, and control systems. The software System IDentification Programs for AirCRAFT (SIDPAC) [12], which is associated with Ref. [13], was used for the input design, equation error and output error parameter estimation, and nonlinear aircraft simulation.

II. Adaptive Tuning Method

II.A. Overview

A flow chart for the adaptive tuning algorithm is shown in Fig. 2. To start, the aircraft is trimmed to a reference condition. The aircraft may be flying open loop or with control laws active. Each control gain is then sequentially perturbed and a dynamic maneuver is conducted to identify closed-loop dynamic models and to compute performance metrics. Afterwards, the local cost gradient is calculated and control gains are updated to improve the performance. This process repeats until convergence to the desired performance.

The following subsections discuss these steps in more detail. Because the method is demonstrated later using a longitudinal SAS, this presentation is specialized for that problem. However, the algorithm is general and can be applied to other aircraft and axes, or to different types of control systems including autopilots, control augmentation systems, or adaptive controllers.

The approach is similar to extremum control or peak-seeking control [14, 15], which originated in the 1940’s and has been successfully applied in recent aerospace applications. For example, peak-seeking control was used to reduce drag of an aircraft in formation flight [16, 17] and to optimize F-18 trim settings

for minimum fuel consumption [18, 19]. For rapid adaptation, sinusoidally varying parameters are used to continually change and identify models before closed-loop dynamics settle [15]. The analogy for this problem would be to use sinusoidally varying control gains, instead of sequential step changes. However, the identification techniques used in this paper are typically formulated for systems with constant parameters. After some experimentation, it was found that more accurate parameter estimates and faster convergence of the control gains was obtained using discretely changing control gains. Using continuously changing control gains required longer durations of time so that accurate dynamic models could be identified before the gains significantly changed.

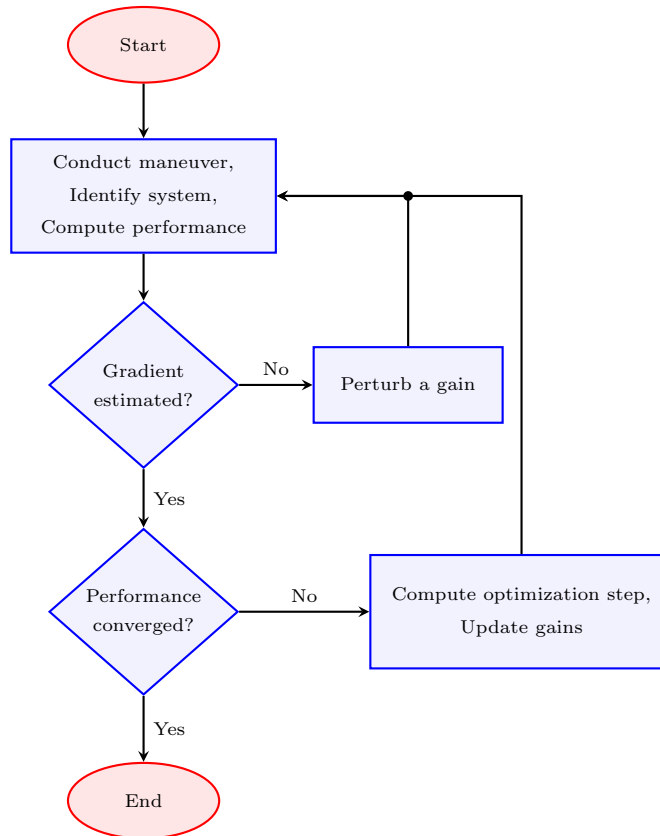


Figure 2. Algorithm flow chart for adaptive tuning of control gains

II.B. LOES Model Structure

A dynamic model of the closed-loop flight dynamics is needed to compute performance metrics, such as natural frequencies or settling times. Although it has been shown during the Learn-to-Fly project that the form and structure of parametric models can be learned onboard the aircraft in real time from measured data [1, 7], most airplanes for the envisioned applications generally behave similarly to classical flight dynamics models. In this case, a Low-Order Equivalent System (LOES) model [20, 21] is appropriate for identification of the closed-loop dynamics.

The block diagram depicted in Fig. 3 shows that the LOES model subsumes all the dynamics between the pilot input and the response of the aircraft, including the bare airframe, flight control system, sensors, and actuators. For the short-period approximation investigated later, the LOES model is a second-order transfer function [21] which can take the state-space representation

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

The states for this model are angle of attack α and pitch rate q , whereas the input is the pilot longitudinal

stick displacement η . In this linear model, the states and control are perturbations from reference values. The equivalent time delay τ approximates latencies due to actuator lags, the control system, nonlinearities, and other factors. The equivalent dimensional stability and control derivatives (e.g., L_α) do not describe the aerodynamic forces and moments, as is the usual case, but the lumped dynamics of the closed-loop system.

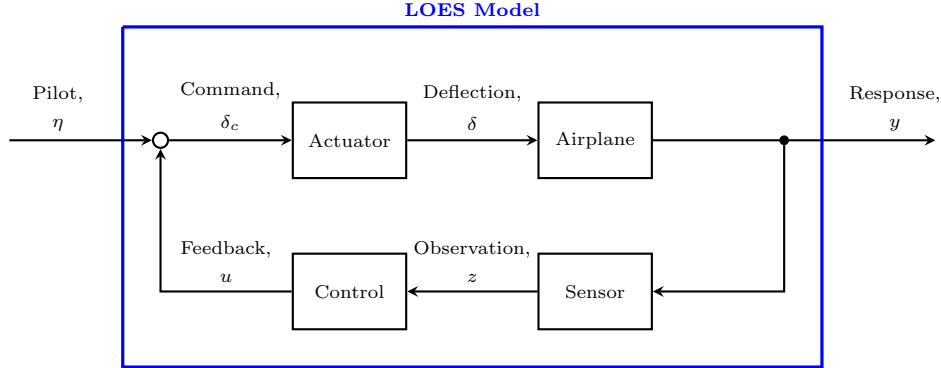


Figure 3. Block diagram of the LOES model

The control block contains the flight control system and the control laws, which operate on sensor outputs and produce feedback to be added to the pilot inputs. The architecture of the control laws is not needed for the adaptive tuning method; only authority over the set of gains \mathbf{k} . For the SAS tuning example discussed later, the gains consisted of k_α and k_q , which are proportional to angle of attack and pitch rate. In the identified LOES models, the effect of these control gains are subsumed into the equivalent stability and control derivatives.

II.C. LOES Parameter Estimation

Parameter estimation was used to determine values and uncertainties for the LOES model terms, which vary with flight condition and control gains. To obtain accurate estimates, the pilot must excite the aircraft to a small but sufficient degree and over a specified bandwidth. This creates small perturbation responses with good information content for the relevant LOES model parameters.

The time delay was first estimated. This parameter is correlated with the stability and control derivatives, but all these parameters can be identified separately and accurately using other information [20]. If both the pilot stick and the control surfaces δ are instrumented, the time delay can be identified in the frequency domain using the model [22]

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

where, for example, $\eta(j\omega)$ denotes the Fourier transform of $\eta(t)$. This is a nonlinear estimation problem that can be solved using output error [13, 23], which is closely related to the gain optimization discussed next in Section II.D. Nonlinear optimizations require iteration and good starting values. However, this estimation of the time delay is well conditioned and runs quickly, and starting values can be initialized to zero. Therefore the nonlinear optimization poses no significant problem during flight testing. Furthermore, this estimation need only be performed once during the tuning procedure because the equivalent time delay characterizes latencies that are not expected to significantly change between maneuvers. If this is not the case, the gain optimization has better convergence if the time delay is updated periodically rather than with every maneuver.

With the time delay determined and fixed, the equivalent stability and control derivatives were next estimated using equation error in the frequency domain [13]. The model equations used in this paper are more conveniently posed in the time domain as

$$-\frac{mg}{\bar{q}S} a_z = C_{L_V} \frac{\Delta V}{V_0} + C_{L_\alpha} \Delta\alpha + C_{L_q} \frac{q\bar{c}}{2V} + C_{L_\eta} \Delta\eta(t - \hat{\tau}) \quad (3a)$$

$$\dot{q} = C_{m_V} \frac{\Delta V}{V_0} + C_{m_\alpha} \Delta\alpha + C_{m_q} \frac{q\bar{c}}{2V} + C_{m_\eta} \Delta\eta(t - \hat{\tau}) \quad (3b)$$

The left side of Eq. (3a) is an approximation for low angles of attack, but the full form can be used if necessary. Although dimensional derivatives are typically used for LOES identification in handling qualities analyses, it was found that estimating nondimensional quantities

$$L_V = \frac{\bar{q}S}{mV} \frac{1}{V_0} C_{L_V} \quad L_\alpha = \frac{\bar{q}S}{mV} C_{L_\alpha} \quad L_q = \frac{\bar{q}S}{mV} \frac{\bar{c}}{2V} C_{L_q} \quad L_\eta = \frac{\bar{q}S}{mV} C_{L_\eta} \quad (4a)$$

$$M_V = \frac{\bar{q}S}{mV} \frac{1}{V_0} C_{m_V} \quad M_\alpha = \frac{\bar{q}S}{mV} C_{m_\alpha} \quad M_q = \frac{\bar{q}S}{mV} \frac{\bar{c}}{2V} C_{m_q} \quad M_\eta = \frac{\bar{q}S}{mV} C_{m_\eta} \quad (4b)$$

had better accuracy because the known functional dependencies on airspeed V and dynamic pressure \bar{q} were accounted for. Airspeed derivatives were not included in the short period LOES model in Eq. (1) but were estimated to keep airspeed variations from being attributed to other parameters. Not including the airspeed derivatives resulted in biased parameter estimates and poor convergence of the gain estimator. Lastly, the vertical accelerometer output a_z was used for Eq. (3a) instead of differentiating the angle of attack measurement because this equation has higher sensitivity for the model parameters and produces more accurate estimates [13].

Each of Eqs. (3) is a linear estimation problem because $\hat{\tau}$ is known and fixed. The equations and data were transformed into the frequency domain and represented as

$$\mathbf{z} = \mathbf{X}\boldsymbol{\theta} + \mathbf{v} \quad (5)$$

where \mathbf{z} is the modeled variable, \mathbf{X} is the regressor matrix of explanatory variables, $\boldsymbol{\theta}$ contains the unknown parameters, and \mathbf{v} is the equation error. The ordinary least-squares solution is

$$\hat{\boldsymbol{\theta}} = [\Re(\mathbf{X}^\dagger \mathbf{X})]^{-1} \Re(\mathbf{X}^\dagger \mathbf{z}) \quad (6)$$

which has the covariance

$$\text{cov}(\hat{\boldsymbol{\theta}}) = \sigma^2 [\Re(\mathbf{X}^\dagger \mathbf{X})]^{-1} \quad (7)$$

The error variance σ^2 is typically computed from the model residuals.

II.D. Control Gain Optimization

Each time a gain is perturbed and a LOES model is identified, the performance metrics are computed and the cost function is evaluated. For the longitudinal SAS example discussed later, the performance metrics

$$\boldsymbol{\gamma} = \begin{bmatrix} \omega_n \\ \zeta \end{bmatrix} \quad (8)$$

were the frequency ω_n and the damping ratio ζ of the short period mode. These could have also been chosen as the real and imaginary parts of the short period pole, or could have included settling times, or could have included phugoid characteristics. The cost function

$$J(\mathbf{k}) = \frac{1}{2} \boldsymbol{\nu}^T \mathbf{S}_{\nu\nu}^{-1} \boldsymbol{\nu} \quad (9)$$

was minimized to determine the optimal feedback control gains

$$\mathbf{k} = \begin{bmatrix} k_\alpha \\ k_q \end{bmatrix} \quad (10)$$

where the performance residuals in the cost function

$$\boldsymbol{\nu} = \boldsymbol{\gamma}_d - \hat{\boldsymbol{\gamma}} \quad (11)$$

are the differences between the desired performance $\boldsymbol{\gamma}_d$ and the performance computed from the identified LOES models $\hat{\boldsymbol{\gamma}}$.

The residual covariance matrix $\mathbf{S}_{\nu\nu}$ describes the uncertainty on the estimated performance metrics and was used in the cost function to prioritize more accurate data over less accurate data. This covariance is

determined by employing the unscented transform [24] to propagate uncertainties on the identified LOES parameters to uncertainties on the computed performance metrics. Analysis using Monte Carlo simulation confirmed this method produced accurate residual covariance estimates in an efficient manner.

After each gain is sequentially perturbed about nominal settings and performance metrics are calculated from the corresponding identified LOES models, the performance sensitivity matrix

$$\frac{\partial \boldsymbol{\gamma}}{\partial \mathbf{k}} = \begin{bmatrix} \frac{\partial \omega_n}{\partial k_\alpha} & \frac{\partial \omega_n}{\partial k_q} \\ \frac{\partial \zeta}{\partial k_\alpha} & \frac{\partial \zeta}{\partial k_q} \end{bmatrix} \quad (12)$$

is computed using finite differences. This matrix quantifies changes in short period frequency and damping ratio from changes in the control gains. Forward or backward finite differences require fewer maneuvers to compute cost gradients, but gradients are estimated more accurately using central finite differences and may lead to fewer iterations overall. In this context, backward finite differences are preferred over forward finite differences because gains are reduced, which safeguards against incurring large responses or instabilities. If the form of the control architecture is known, this matrix could be computed by substituting measured data into its analytical expression to perhaps reduce the number of maneuvers needed for convergence.

Control gain perturbations should be small enough so that computed gradients approximate tangent planes to the cost surface rather than secant planes. However, perturbations should be large enough so that a change in the performance can be distinguished from the effects of noise and modeling error. For this paper, gain perturbation sizes were adjusted for each iteration step so that the predicted pole locations of the identified LOES model would move by 10% in the complex plane. In flight, a pilot or test engineer could also select gain perturbations by observing the aircraft responses.

After all the gains have been perturbed, the local cost gradient

$$\frac{\partial J}{\partial \mathbf{k}} = -\frac{\partial \boldsymbol{\gamma}^T}{\partial \mathbf{k}} \hat{\mathbf{S}}_{\nu\nu}^{-1} \boldsymbol{\nu} \quad (13)$$

and information matrix (which is an approximation of cost function Hessian)

$$\mathbf{M} = \frac{\partial \boldsymbol{\gamma}^T}{\partial \mathbf{k}} \hat{\mathbf{S}}_{\nu\nu}^{-1} \frac{\partial \boldsymbol{\gamma}}{\partial \mathbf{k}} \quad (14)$$

are computed. The control gains are then updated for the next iteration as

$$\hat{\mathbf{k}}_i = \hat{\mathbf{k}}_{i-1} - \mathbf{M}^{-1} \frac{\partial J}{\partial \mathbf{k}} \quad (15)$$

and the uncertainties for these estimates are

$$\text{cov}(\hat{\mathbf{k}}) = \mathbf{M}^{-1} \quad (16)$$

These steps are called a Gauss-Newton optimization, which is a modification of the widely-used Newton-Raphson method. The advantage of using this optimization is that second derivatives of the cost function, which would require more maneuvers to compute, are not needed. To start this optimization, pole-placement techniques can be used to obtain starting values for the control gains near a solution. However, the solution from pole placement is expected to be approximate because it uses an identified linear approximation of the dynamics using different control gain settings.

After the gains are updated using a Gauss-Newton step, the process of perturbing the gains repeats until convergence to the desired performance is achieved. There are many metrics for determining convergence. For example, the estimated performance may fall within a specified tolerance of the desired performance, or convergence could be defined as when values of cost, cost gradient, parameters, and uncertainties vary within a tolerance between maneuvers.

III. Nonlinear Aircraft Simulation

The Free-flying Aircraft for Sub-scale Experimental Research (FASER) project aimed to provide an inexpensive and subscale airplane testbed for flight research into experiment design, data analysis, dynamic modeling, and control law design.²⁵ The test aircraft for this project was the Hanger 9 Ultra-Stick™ 120 kit-built tail-dragger pictured in Fig. 4. The airplane is a subscale, high-wing, general aviation-type aircraft with an electrically-driven propeller. Nominal mass properties and geometry information for the aircraft are listed in Table 1.



Figure 4. FASER test aircraft (source: Ref. [26])

Table 1. Test airplane geometry and nominal mass properties

Parameter	Value	Unit
b	6.33	ft
\bar{c}	1.45	ft
S	9.18	ft ²
x_{cm}/\bar{c}	0.24	–
m	0.51	slug
I_{xx}	1.32	slug-ft ²
I_{yy}	1.95	slug-ft ²
I_{zz}	0.63	slug-ft ²
I_{xz}	0.32	slug-ft ²

The nonlinear flight dynamics simulation for this aircraft is described in Ref. [11]. This simulation was chosen for analysis because the aircraft was similar in scale to other subscale aircraft used during the Learn-to-Fly project. The aerodynamic database for this simulation uses multivariate polynomial splines generated from wind tunnel and first-principles analyses, tabulated between -10 to $+30$ degrees angle of attack and 0.05 to 0.3 Mach. Conventional throttle, elevator, aileron, and rudder controls are modeled for throttle settings $0-1$, elevator deflections between ± 25 degrees, aileron deflections between ± 25 degrees, and rudder deflections between ± 30 degrees. An engine model includes first-order lag dynamics from the throttle input to engine power, and is valid for thrust between $0-10$ lbf.

The nonlinear simulation discussed in Ref. [11] was modified for this work. First, actuator models

$$\frac{\delta(s)}{\delta_c(s)} = \frac{18.8}{s + 18.8} \quad (17)$$

were added for each control surface. These models added realism to the simulation but also made values of the LOES model parameters different than those for the open-loop bare airframe. Second, the simulation was reconfigured such that inputs were pilot commands rather than control surface deflections. This change allowed the longitudinal SAS to be implemented using the elevator command

$$\delta_c = \eta - k_\alpha \Delta\alpha - k_q \Delta q \quad (18)$$

which adds proportional feedback from α and q to the pilot command. Third, sensor noise was added to the simulated measurements. The noise standard deviations used given in Table 2, and were obtained from flight test data with other subscale aircraft. Lastly, all pilot inputs were passed through the filter

$$H(s) = \frac{12.57}{s + 12.57} e^{-0.1s} \quad (19)$$

to approximate the dynamics of a human pilot.

Table 2. Sensor noise values

Sensor output	Std. dev.	Unit
η	0.010	deg
δ	0.025	deg
V	0.092	ft/s
α	0.082	deg
q	0.234	deg/s
a_z	0.004	g

IV. Simulation Results

The open-loop, nonlinear simulation was trimmed for steady, wings-level flight at the nominal reference condition listed in Table 3. The bare airframe dynamics obtained from numerical linearization of the original FASER simulation had short period frequency 7.57 rad/s (1.20 Hz) and damping ratio 0.75. Applying the same procedure to the augmented FASER simulation and residualizing the actuator dynamics lowered the short period frequency to 7.29 rad/s (1.16 Hz).

Table 3. Reference condition for FASER nonlinear simulation

Parameter	Value	Unit
V	80.0	ft/s
α	2.46	deg
δ_t	0.24	–
δ	–2.01	deg
h	1000	ft

The goal for tuning the SAS feedback gains k_α and k_q was to achieve a closed-loop short period frequency $\omega_n = 5.03$ rad/s (0.8 Hz) and damping ratio $\zeta = 0.7$, within 2% of each of these parameters. The purpose was to slow the short period response and make the aircraft easier to fly by a remote pilot, who might be contending with time delays, visibility issues, and other factors during flight testing. These adjustments to the short period dynamics were relatively minor and only small control gains were expected.

To help understand the tuning results presented next, a search was conducted to interrogate the cost surface and find by trial and error the optimal control gains to achieve the desired closed-loop performance of the LOES model. This was performed using the same excitation input and procedure described below, but was done without measurement noise on the output data to gain clearer insight. The resulting cost surface was non-convex and contained several local minima. Large peaks in the cost surface arose when the control gains produced unstable modes or two first-order poles rather than a second-order pole. Figure 5 shows a portion of the cost surface near the optimal solution. To the resolution shown, the optimal gains were $k_\alpha = -0.05$ and $k_q = 0.08$, which is plotted as the red dot. The cost surface contained a valley where many

combinations of gains produce nearly the same cost. This is because both gains affected both performance metrics in this region. For example, in the open-loop case, the gain sensitivities were

$$\frac{\partial \gamma}{\partial \mathbf{k}} = \begin{bmatrix} \frac{\partial \omega_n}{\partial k_\alpha} & \frac{\partial \omega_n}{\partial k_q} \\ \frac{\partial \zeta}{\partial k_\alpha} & \frac{\partial \zeta}{\partial k_q} \end{bmatrix} = \begin{bmatrix} -5.28 & -46.8 \\ 1.03 & 1.65 \end{bmatrix}$$

The damping ratio was almost equally sensitive to the two gains, whereas ω_n was approximately 9 times more sensitive to k_q than k_α . Due to this sensitivity, the valley in the cost function was smaller in the k_q dimension than in k_α . Within the valley was also a deeper trench, along which were the optimal gains.




Figure 5. Cost surface for gain optimization

The pilot excitation input was chosen to be a multisine input, which contained seven phase-optimized sinusoids over the range 0.3–2.1 Hz. The measured input and responses are shown for the first maneuver of the tuning procedure in Fig. 6. The pilot and actuator filters attenuated some of the higher-frequency content in the input, and the responses were small perturbations. Although multisines are advanced inputs which are typically implemented by the flight control software, pilots can approximate this input reasonably well [27]. The important aspects of this excitation for this application are only that it should contain sufficient excitation over a bandwidth enclosing the short period mode, and that it should keep responses about the reference condition small for accurate linear modeling. In addition to the multisine, frequency sweeps and 3-2-1-1 multistep inputs also led to good results using the simulation data. Pulses, doublets, and 2-1-1 multistep inputs, however, did not have enough data information, and inaccuracies in the identified LOES model led to poor convergence or divergence in the gain estimation.

Data from the first maneuver were used to estimate the equivalent time delay and LOES model. The frequency range for the analysis was 0.17–2.5 Hz, which encompassed the excitation and short period response. Good model fits to frequency-domain data for the first excitation maneuver are shown in Fig. 7. The time delay was estimated to be 0.0491 ± 0.0012 . The standard error was approximately 2.4% of the parameter value, which is small. This time delay is due to the actuator dynamics, which impacts the estimation of all other parameters. The time delay value was fixed for all subsequent maneuvers, as this was not expected to change and doing so improved the convergence of the gain estimation.

The estimated stability and control derivatives also had good fits to the frequency-domain data, as shown in Fig. 7, with low standard errors. Most explanatory variables had low pairwise correlations. However, α and q had correlation 0.84, whereas q and $\eta(t - \hat{\tau})$ had correlation -0.80 . These correlation magnitudes were high but remained below the guideline of less than 0.9 in absolute value for modeling [13]. The identified LOES model had short period frequency 7.06 rad/s (1.12 Hz) and damping ratio 0.77, which is similar to the true values for the system mentioned earlier.

Pole placement was then employed to compute gains near a solution to improve the convergence of the gain optimization. With the identified LOES model, pole placement predicted gains $k_\alpha = -0.0259$ and $k_q = 0.0659$ to achieve the desired performance. As seen from the cost surface shown in Fig. 5, these values




Figure 6. Measurements from the first maneuver




Figure 7. Fits to frequency-domain data for estimating LOES model parameters

were close to the optimized values for the control gains. The subsequent LOES model identified from the second excitation maneuver and using these gains had the short period frequency 4.72 (0.75 Hz) and damping ratio $\zeta = 0.76$. These values were within 6.1% and 8.6% of the desired performance metrics, but did not meet the 2% criterion. The nonlinear gain optimization was then used to further refine the control gain estimates.

Only one iteration of the gains using the nonlinear optimization was needed to achieve the desired frequency and damping ratio within the specified 2% tolerance. Backwards finite differences provided enough accuracy for the cost gradient computations and reduced the number of maneuvers that would have been performed using central finite differences. Five maneuvers were therefore used: one with the aircraft in the initial open-loop configuration, one after setting the gains using pole placement, and three to adjust and update the control gains using the gain optimization. These five maneuvers each began at the reference condition and are shown as concatenated time histories in Fig. 8. Numbers above the plot indicate the maneuver number. Although just over a minute of test time was needed to achieve the optimal gain estimates, additional time would be needed in flight to change the control gains and retrim the aircraft between maneuvers. In flight, a pilot would change the stick inputs to keep responses at roughly the same amplitude. Because the amplitude variations were small and remained in the linear range, the stick input amplitude was not altered between maneuvers. However, the control gain adjustments were large enough that the effect can be seen in the elevator and response data.




Figure 8. Measurements for all five excitation maneuvers

Estimates of the equivalent stability and control derivatives and the associated 2σ uncertainties are shown for the five maneuvers in Fig. 9. From analysis of the first maneuver, the time delay was fixed at 0.0491. The airspeed derivatives are not shown because they were only included to obtain accurate estimates of the other parameters and were not included in the LOES model. The derivatives L_α , L_q , L_η , and M_α did not vary much with the changes in gain. However, M_η had a significant shift when the gains were first adjusted, and M_q became statistically insignificant. This was due to an increase in the correlation between α and q mentioned earlier.

Poles for the LOES models are plotted (for one of the two conjugate pairs) on the complex plane in Fig. 10. The number appearing next to each pole indicates the maneuver number. The desired pole is shown in red, where the shaded region encompasses the 2% error range in frequency and damping ratio. The first maneuver resulted in the pole on the far left side in the plot. The second pole, which corresponded

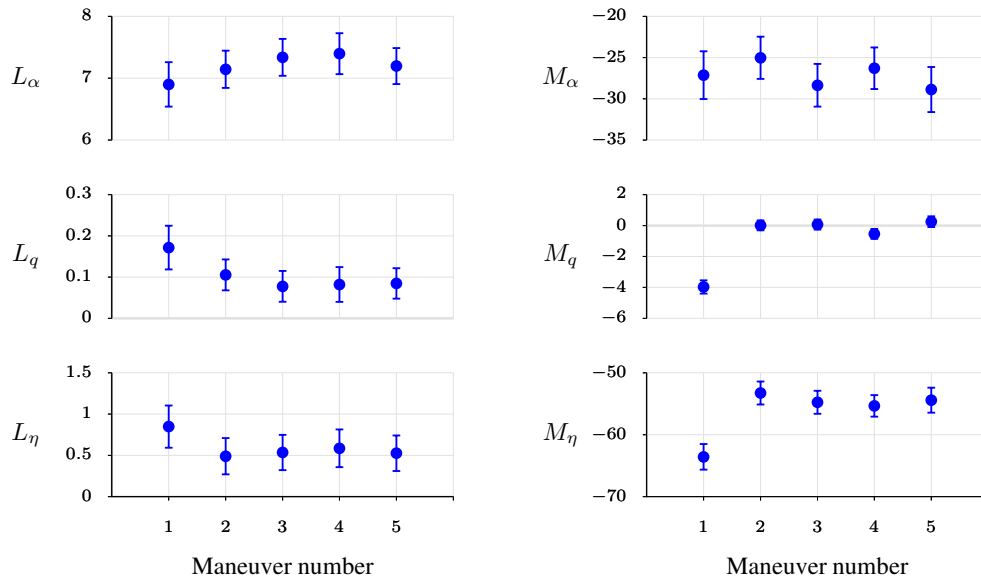


Figure 9. Equivalent stability and control derivative estimates

to gains predicted using pole placement, was much closer to the desired performance. The subsequent two poles above and left are from LOES parameters identified from maneuvers where k_α and k_q were perturbed, respectively. From these three poles, the local cost gradient was estimated and the resulting Gauss-Newton update of the gains led to the final pole, which was within 2% of the desired performance.

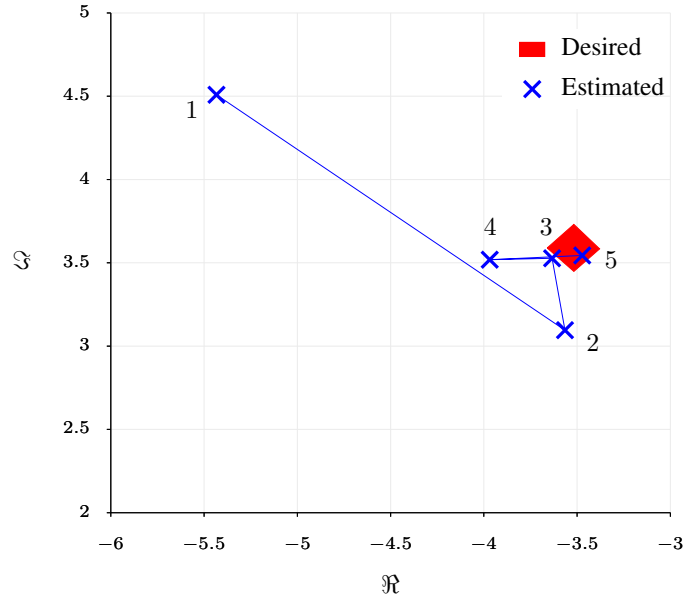


Figure 10. Estimated and desired poles for the identified LOES models

Figure 11 shows the short period frequency and damping ratios. The 2σ uncertainties were computed from propagating uncertainties in the equivalent stability and control derivatives to the performance metrics using the unscented transform. After pole placement, the frequency was within statistical range of the desired value, and only a few more maneuvers were needed to place its value within the 2% tolerance. The change in damping ratio is small, and has larger uncertainty. By the fifth maneuver, both performance metrics had less than 1.3% error, which was within the 2% specification.

The estimated control gains are shown in Fig. 12 with 2σ uncertainties. These gains were approximately




Figure 11. Estimated performance metrics

the same as those shown in the cost surface in Fig. 5. The k_q gains were larger and less uncertain than the k_α gains because they had stronger sensitivities. This is consistent with the shorter valley in the k_q dimension in Fig. 5.




Figure 12. Estimated control gains

To check that the nonlinear simulation behaved as expected with the tuned SAS controller, step responses were examined and are shown in Fig. 13. The input was a doublet intended to excite the short period mode. The responses in blue are the open-loop nonlinear simulation responses, before the tuning. The red lines are the desired responses from the identified LOES model using the last maneuver and the optimized control gains. The green lines are the responses from the nonlinear simulation using the optimized control gains. The closed-loop nonlinear responses matched the predicted responses well. The small discrepancies could be from neglecting the airspeed dynamics. Regardless, this verified that the closed-loop nonlinear simulation behaved as expected.

The simulation results here showed one flight test in which five maneuvers were needed for convergence. Observing repeated simulation runs, the tuning procedure often converged in five maneuvers. It was also common that a second optimization step was needed, which resulted in 8 total maneuvers. In rare cases, only three iterations were needed if the gains were perturbed in a fortuitous manner during the first optimization




Figure 13. Comparison of doublet responses before and after gain tuning

step. To reduce the number of maneuvers, the gain perturbations could be reordered for each optimization step to increase the likelihood of this occurring.

Similar numbers of maneuvers were needed when using forward finite differences instead of backward finite differences. Using central finite differences, the optimization consistently converged using one optimization step. However, because each control gain was perturbed twice, this resulted in 7 maneuvers. The cost of computing more accurate gradients was that more data and more maneuvers were needed. For systems with a small number of gains, it may be more efficient to use central finite differences, rather than backward finite differences, and have a higher chance of converging within one optimization step. This would also help safeguard against effects of noise, modeling error, and low information content. This could, however, become prohibitive if a large number of control gains were to be tuned.

The number of iterations also depend on other factors, such as noise and turbulence levels. Frequency domain analysis is generally robust to measurement noise, and the equation-error approach considers process noise (turbulence); however, increased levels typically resulted in additional numbers of maneuvers. More analysis and flight testing is needed, but these simulation results indicate that this adaptive tuning method is feasible for tuning control gains in a flight control system during flight.

V. Conclusions

A 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 low-order equivalent system models. Cost gradients are then computed and an online nonlinear optimization is used to update the nominal control gain values to improve performance. This process is repeated during flight until the closed-loop performance is within a specified tolerance of the desired performance.

A simulated experiment was presented using the FASER subscale aircraft nonlinear simulation. Results indicated that the method performed well using a multisine, frequency sweep, and 3-2-1-1 excitation inputs, although only results using the multisine input were shown. Backward differences were found to be sufficient for good convergence and reduced the number of excitation maneuvers needed. LOES parameter estimates were found to be accurate, but some correlations in the data produced questionable values of two pitching moment derivatives. Pole placement was used after the first optimization iteration and put the performance within 8.6% of the desired performance. This estimate was then refined using the gain optimization and three maneuvers were used to reduce this error to 1.3%. Using the optimal gain predictions, the closed-loop nonlinear simulation responded to doublet inputs as expected. From these simulation results, this adaptive tuning method appears feasible for real-time tuning of aircraft control gains during flight.

Acknowledgments

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

References

- ¹Morelli, E., “Real-Time Global Nonlinear Aerodynamic Modeling for Learn-To-Fly,” No. 2016-2010 in Atmospheric Flight Mechanics Conference, AIAA, San Diego, CA, January 2016.
- ²Grauer, J., “Aircraft Fault Detection using Real-Time Frequency Response Estimation,” No. 2016-0372 in Guidance, Navigation, and Control Conference, AIAA, San Diego, CA, January 2016.
- ³Mekky, A. and Gonzalez, O., “LQ Control for the NASA Learn-to-Fly Free-to-Roll Project,” National Aerospace and Electronics Conference, IEEE, Dayton, OH, July 2016, pp. 173–178.
- ⁴Murphy, P. and Brandon, J., “Efficient Testing Combining Design of Experiment and Learn-to-Fly Strategies,” No. 2017-0696 in Atmospheric Flight Mechanics Conference, AIAA, Grapevine, TX, January 2017.
- ⁵Heim, E., Viken, E., Brandon, J., and Croom, M., “NASA’s Learn-to-Fly Project Overview,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ⁶Riddick, S., Busan, R., Cox, D., and Laughter, S., “Learn-to-Fly Test Setup and Concept of Operations,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ⁷Morelli, E., “Practical Aspects of Real-Time Modeling for the Learn-to-Fly Concept,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ⁸Foster, J., “Autonomous Guidance Algorithms for NASA Learn-to-Fly Technology Development,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ⁹Snyder, S., Bacon, B., Morelli, E., Frost, S., Teubert, C., and Okolo, W., “Online Control Design for Learn-to-Fly,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ¹⁰Weinstein, R., Hubbard, J., and Cunningham, M., “Fuzzy Modeling and Parallel Distributed Compensation for Aircraft Flight Control from Simulated Flight Data,” Atmospheric Flight Mechanics Conference, AIAA, Atlanta, GA, June 2018, (to be published).
- ¹¹Garza, F. R. and Morelli, E. A., “A Collection of Nonlinear Aircraft Simulations in MATLAB,” Tech. Rep. TM-2003-212145, NASA, Hampton, VA, January 2003.
- ¹²Morelli, E., “System IDentification Programs for AirCRAFT (SIDPAC),” <http://software.nasa.gov>, Accessed: 2018-05-16.
- ¹³Morelli, E. A. and Klein, V., *Aircraft System Identification: Theory and Practice*, Sunflyte, Williamsburg, VA, 2nd ed., 2016.
- ¹⁴Åström, K. A. and Wittenmark, B., *Adaptive Control*, Series in Electrical Engineering: Control Engineering, Addison Wesley, Reading, MA, 2nd ed., 1995.
- ¹⁵Ariyur, K. B. and Krstić, M., *Real-Time Optimization by Extrimum-Seeking Control*, Wiley, Hoboken, NJ, 2003.
- ¹⁶Ryan, J. J. and Speyer, J. L., “Peak-Seeking Control Using Gradient and Hessian Estimates,” No. WeA17.5, American Control Conference, Baltimore, MD, June 2010.
- ¹⁷Hanson, C. E., Ryan, J. J., Allen, M. J., and Jacobson, S. R., “An Overview of Flight Test Results for a Formation Flight Autopilot,” No. 2002-4755 in Guidance, Navigation, and Control Conference, AIAA, Monterey, CA, January 2002.
- ¹⁸Schaefer, J. and Brown, N. A., “Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Architecture and Performance Metrics,” No. 2013-4856 in Guidance, Navigation, and Control Conference, AIAA, Boston, MA, August 2013.
- ¹⁹Brown, N. A. and Schaefer, J., “Peak-Seeking Optimization of Trim for Reduced Fuel Consumption: Flight-Test Results,” No. 2013-5171 in Guidance, Navigation, and Control Conference, AIAA, Boston, MA, August 2013.
- ²⁰Morelli, E. A., “Identification of Low Order Equivalent Systems Models from Flight Test Data,” Tech. Rep. 2000-210117, NASA, Hampton, VA, August 2000.
- ²¹Anon., “Flying Qualities of Piloted Aircraft,” Tech. Rep. Mil-Std 1797a, Department of Defence, 1990.
- ²²Morelli, E. A., “Dynamic Modeling from Flight Data with Unknown Time Skews,” *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 8, August 2017, pp. 2083–2091.
- ²³Maine, R. E. and Iliff, K. W., “Application of Parameter Estimation to Aircraft Stability and Control: The Output-Error Approach,” Tech. Rep. RP-1168, NASA, Edwards, CA, June 1986.
- ²⁴Uhlmann, J., *Dynamic Map Building and Localization: New Theoretical Foundations*, Ph.D. thesis, University of Oxford, 1995.
- ²⁵Owens, D. B., Cox, D. E., and Morelli, E. A., “Development of a Low-Cost Sub-Scale Aircraft for Flight Research: The FASER Project,” No. 2006-3306 in Aerodynamic Measurement Technology and Ground Testing Conference, AIAA, San Francisco, June 2006.
- ²⁶Denham, C. and Owens, D. B., “Rotary Balance Wind Tunnel Testing for the FASER Flight Research Aircraft,” No. 2016-3105 in Atmospheric Flight Mechanics Conference, AIAA, Washington, D.C., June 2016.
- ²⁷Brandon, J. and Morelli, E., “Real-Time Onboard Global Nonlinear Aerodynamic Modeling from Flight Data,” *Journal of Aircraft*, Vol. 53, No. 5, September–October 2016, pp. 1261–1297.