

\mathcal{L}_1 Adaptive Control with Switched Reference Models: Application to Learn-to-Fly

Steven Snyder *

NASA Langley Research Center, Hampton, VA 23681

Pan Zhao[‡] and Naira Hovakimyan[§]

University of Illinois at Urbana-Champaign, Urbana, Illinois 61801

Learn-to-Fly (L2F) is a new framework that aims to replace the traditional iterative development paradigm for aerial vehicles with a combination of real-time aerodynamic modeling, guidance, and learning control. To ensure safe learning of the vehicle dynamics on the fly, this paper presents an \mathcal{L}_1 adaptive control (\mathcal{L}_1 AC) based scheme, which actively estimates and compensates for the discrepancy between the intermediately learned dynamics and the actual dynamics. First, to incorporate the periodic update of the learned model within the L2F framework, this paper extends the \mathcal{L}_1 AC architecture to handle a switched reference system subject to unknown time-varying parameters and disturbances. The paper also includes analysis of both transient and steady-state performance of the \mathcal{L}_1 AC architecture in the presence of non-zero initialization error for the state predictor. Second, the paper presents how the proposed \mathcal{L}_1 AC scheme is integrated into the L2F framework, including its interaction with the baseline controller and the real-time modeling module. Finally, flight tests on an unmanned aerial vehicle (UAV) validate the efficacy of the proposed control scheme.

I. Introduction

Historically, aerodynamic modeling from computational fluid dynamics (CFD) data, wind tunnel test data, or flight test data has been a multi-step process where the tests are performed, followed by post-test analysis of the data. During this data analysis, it is often discovered that the data content has insufficient richness in particular areas, requiring additional testing to fill in the data gaps. This results in a time-consuming process that may involve a large number of test sorties to acquire the necessary data for development of an adequate aerodynamic model. In view of these disadvantages, there has been an effort to develop the so-called *Learn-to-Fly* (L2F) framework, which aims to replace the current iterative vehicle development paradigm with a combination of real-time aerodynamic modeling, learning control, and other techniques, enabling real-time, self-learning flight vehicles [1]. [The advantages of inflight modeling](#)

*Research Engineer, Dynamics Systems and Control Branch; steven.m.snyder@nasa.gov. Member AIAA.

[‡]Postdoctoral Researcher, Department of Mechanical Science and Engineering; panzhao2@illinois.edu. Corresponding author.

[§]Professor, Department of Mechanical Science and Engineering; nhovakim@illinois.edu. Fellow AIAA.

include the ability to analyze model fidelity and data richness while gathering data, which can in turn reduce the number of test flight sorties required. The conventional aircraft development process and the L2F concept are depicted in Figure 1. Within the L2F framework, state-of-the-art in-flight aerodynamic modeling approaches are combined with guidance and learning control law design methods on board the aircraft, running in real time. However, during the initial learning phase, the only available model of the vehicle’s aerodynamics is based on a guess, which may be very inaccurate, depending on the amount of ground testing performed. The control law must be able to ensure stability of the vehicle with this potentially poor model. As the vehicle flies and additional data are collected, the aerodynamic model is improved and, consequently, the controller should adjust to provide a better closed-loop response, e.g., improved tracking performance. In the subsequent narration, we use real-time modeling, model learning and system identification interchangeably.

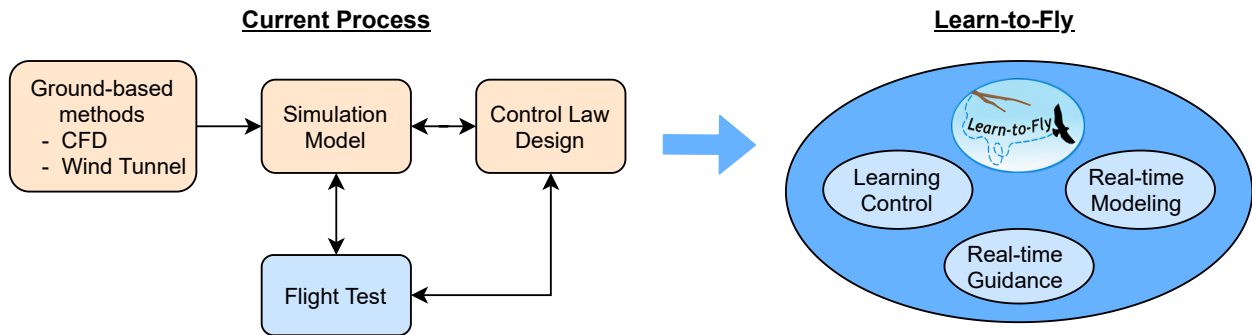


Fig. 1 Conventional aircraft development process vs. Learn-to-Fly concept (reproduction of Fig. 1 in [1])

Since the available dynamic model of the vehicle could be quite poor at the initial stage of model learning, the on-board controller should be capable of stabilizing the vehicle despite the presence of large uncertainties and disturbances. Among different control techniques, adaptive control is one of the most promising methods to handle large uncertainties while providing good nominal performance. In the last several decades, the field of adaptive control has witnessed tremendous developments, capturing different classes of linear and nonlinear systems, subject to various parametric and state-dependent uncertainties, and unmodeled dynamics [2, 3]. As one of the most well-known adaptive control techniques, model reference adaptive control (MRAC) deserves special attention. However, as experience has shown, conventional MRAC controllers lose robustness guarantees in the presence of fast adaptation, and therefore are only effective for slowly-varying uncertainties. \mathcal{L}_1 adaptive control (\mathcal{L}_1 AC) resolved this issue with an appropriate filtering structure that decouples the estimation loop from the control loop [4, 5]. As a result, arbitrarily fast adaptation, subject only to hardware limitations, can be used without sacrificing robustness. Note that fast adaptation is desirable for timely estimation of and compensation for the (fast-varying) uncertainties/disturbances. Benefiting from the fast and robust adaptation, \mathcal{L}_1 AC provides guaranteed tracking performance, both in transient and steady-state, and robustness (e.g., nonzero time delay margin), in the presence of various uncertainties and disturbances [4]. We note that \mathcal{L}_1 AC has

been verified on a large number of real-world systems [4, 6], including NASA’s subscale commercial jet [7, 8], and Calspan’s Learjet (a piloted aircraft) [9, 10].

When incorporating the adaptive control scheme into the L2F framework, we need to address a crucial issue, i.e., updating the reference model as the data collection and model learning processes evolve. For instance, at some point, assuming the real-time modeling module produces a model, M_1 , we may use M_1 to produce a reference model to design an adaptive controller that tries to cancel the discrepancy between the actual uncertain dynamics and M_1 . After a certain amount of time, assuming the model is updated to be M_2 , a natural choice is to use this new model, M_2 to replace the old model, M_1 , when generating the reference model for adaptive control design. Additionally, one might wish to change the specified desired dynamics based on the updated model to ensure model matching conditions are satisfied or to choose a more natural set of desired dynamics from the potentially infinite options that meet the given design criteria (such as those in military or industry standards). Locking in a specific set of desired dynamics when little is known about the vehicle takes away the option of many other viable designs. Essentially, implementation of the L2F framework leads to a switched reference model for an adaptive control scheme. Therefore, we need to address the stability and performance of such an adaptive control scheme. From the above description, it should be clear that learning in the L2F framework is used to update the reference model which characterizes the desired dynamics for the proposed \mathcal{L}_1 AC scheme. On the other hand, in many other adaptive control schemes, e.g., concurrent learning MRAC [11, 12], the desired dynamics are usually pre-specified and fixed, and learning is used for estimation of uncertain parameters that is further used in the control law to cancel the uncertainties

There have been only a few attempts to address adaptive control design for a switched reference model. Sang and Tao studied MRAC design for a piecewise linear (PWL) system with parametric uncertainties [13], in which asymptotic tracking performance is achieved under arbitrarily fast switching when a common Lyapunov function exists for the PWL reference system. In the absence of such common Lyapunov function, only small tracking error in the mean-squared sense can be achieved under slow switching conditions. Xie and Zhao [14] studied MRAC design for a switched linear parameter-varying (LPV) system subject to parametric uncertainties under average-dwell-time switching. Additionally, Wu et al. [15] investigated the issues of asynchronous switching in MRAC design for a switched linear system with parametric uncertainties. All these methods are based on a direct MRAC scheme, and only asymptotic tracking performance can be achieved at most. On the contrary, the \mathcal{L}_1 AC scheme proposed in this paper and also in our recent work [16], to the best of our knowledge, is the first indirect adaptive control scheme considering a switched reference model. Additionally, the \mathcal{L}_1 AC scheme allows for consideration of time-varying parametric uncertainties and (time-dependent) disturbances instead of just the time-invariant parametric uncertainties considered in [13, 14, 16]. Finally, with the help of the \mathcal{L}_1 AC architecture, we will provide guaranteed stability and transient tracking performance in the presence of a switched reference model.

Statement of Contributions. Motivated by the L2F concept, this paper presents an \mathcal{L}_1 AC based scheme to ensure

safe learning of unknown dynamics of an aerial vehicle. The proposed \mathcal{L}_1 AC is intended to provide the necessary robustness required to keep the vehicle flying while the learning module (e.g., [17]) learns the dynamics. At the core is an \mathcal{L}_1 AC architecture with a piecewise constant (PWC) adaptive law for a switched linear system to characterize the change of the desired/nominal dynamics as a result of periodic model update from the real-time system identification module. For stability and performance analysis, we introduce a virtual reference system that depends on the true uncertainties and represents the best achievable performance. Under reasonable assumptions, we derive transient performance bounds in terms of the inputs and states of the adaptive system as compared to the same signals of the reference system. The influence of re-initialization error associated with the state predictor, a key component of an \mathcal{L}_1 controller, is also analyzed. We then illustrate how the proposed \mathcal{L}_1 AC scheme is incorporated into the L2F framework. We finally present the results of flight tests of the proposed control and learning schemes on a fixed-wing unmanned aerial vehicle (UAV) conducted by NASA Langley Research Center (LaRC) in Virginia, USA. To the best of our knowledge, these flight tests were the first of this vehicle class enabled by the L2F concept, i.e., reliance on a combination of real-time modeling and robust adaptive control while minimizing time-consuming and expensive ground testing. To summarize, our contributions include

- (1) an \mathcal{L}_1 AC architecture with a PWC adaptive law for a switched linear system subject to time-varying parametric uncertainties and disturbances, with guaranteed robustness and transient performance;
- (2) an illustration of how the proposed \mathcal{L}_1 AC architecture is incorporated into the control law design and update within L2F framework; and
- (3) validation of the proposed \mathcal{L}_1 AC in the L2F framework with an existing learning scheme [17] using flight tests on a fixed-wing UAV.

In comparison, our preliminary work [16] addresses only time-invariant parametric uncertainties, uses a projection-based adaptive law that is not as amenable to implementation as the PWC law adopted in this paper (more details given in the first paragraph of Section III), is not relevant to L2F, and is only validated on a numerical example. The current work uses L2F as motivation for its development of \mathcal{L}_1 AC with switched desired dynamics, and is also applicable in other situations where it may be advantageous for the desired dynamics to switch, such as having different dynamics for different tasks (e.g., precision control takeoff/landing vs. energy-efficient cruise/loiter) or when vehicle mass properties change (e.g., package pick up and delivery). We do not seek to debate the benefits of the L2F paradigm.

The paper is organized as follows. Section II gives an overview of the L2F framework. Section III presents an \mathcal{L}_1 AC architecture for a switched linear system with stability and performance analysis. Section IV introduces the design of the L2F control system, including both the baseline and adaptive controllers, while Section V includes the test results on a UAV.

The notations are defined as follows. We denote by \mathbb{R}^n the n -dimensional real vector space, and $\mathbb{R}^{m \times n}$ the set of real m by n matrices. Let \mathbb{I}_n denote an n by n identity matrix, and 0 denote the scalar zero or a zero matrix of compatible

dimension. We use $\mathfrak{L}[x(t)]$ to denote the Laplace transform of a signal $x(t)$. We use $\|\cdot\|$ to denote the 2-norm of a vector or matrix. The maximal and minimal singular values of a matrix P are denoted by $\bar{\lambda}(P)$ and $\underline{\lambda}(P)$, respectively. For symmetric matrices P and Q , $P > Q$ means $P - Q$ is positive definite.

II. Overview of the Learn-to-Fly Framework

Learn-to-Fly (L2F) is an advanced technology development effort aimed at merging real-time aerodynamic modeling, learning control, real-time guidance, and other disciplines in order to replace the current iterative vehicle development process with aggressive flight testing of self-learning vehicles (see Fig. 1). Figure 2 presents the L2F framework, and below we provide a brief overview of its primary components. For more details, the reader is directed to [1] and references therein.

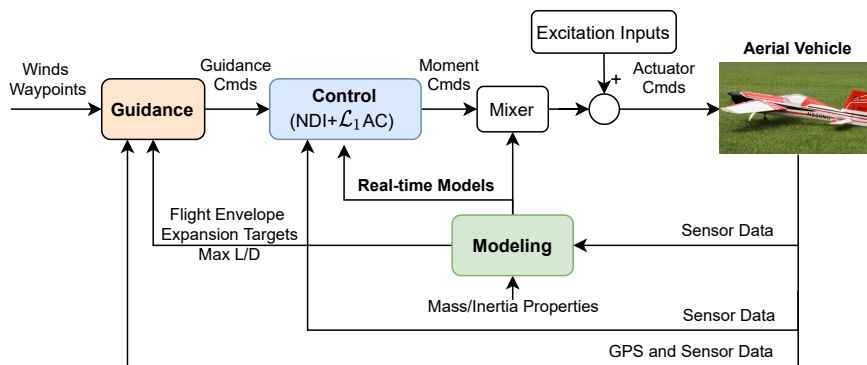


Fig. 2 Block diagram of the Learn-to-Fly architecture (reproduction of Fig. 2 in [1])

Modeling: One of the main objectives of L2F flight tests is to identify a six-degree-of-freedom mathematical model for each non-dimensional aerodynamic force and moment in real time during the flight. The non-dimensional forces and moments must be computed from measured aircraft states, known vehicle geometry, and known mass properties since it is not possible to directly measure these properties during flight [17]. The measured explanatory variables include aircraft states and control surface positions. Two primary factors in developing an accurate global aerodynamics model are the design of efficient flight maneuvers to sufficiently excite the desired aircraft dynamics and the implementation of a real-time recursive system identification scheme [17].

To provide sufficient excitation to the vehicle dynamics, orthogonal phase-optimized multi-sine inputs were injected into the control surface commands during portions of the flight tests [17]. These inputs are called programmed test inputs (PTIs). Because the inputs are orthogonal, the effects of each of the inputs can be discerned despite the fact that multiple inputs are excited simultaneously. Since these inputs are zero-mean and phase-optimized to limit the peak amplitude, their net effect on the aircraft only causes small perturbations from the nominal flight trajectory.

The real-time recursive system identification scheme used was based on multivariate orthogonal functions [17]. This approach uses a pool of candidate multivariate functions, called regressors, of a user-defined level of complexity.

In theory, this pool of candidates could be arbitrarily large, but, in practice, the size of the pool is limited by the onboard computing capability [17]. The explanatory capability unique to each regressor is isolated, quantified, and ranked by orthogonalizing the regressors with a recursive QR decomposition. Coefficients for each orthogonalized regressor are calculated using a least-squares estimator that minimizes the sum of squared differences between the computed non-dimensional forces and moments and those predicted by the model. The model structure is determined based on measures of model fit quality and a penalty for model complexity. These statistical metrics determine which orthogonalized regressors are included in the mathematical model. Finally, the model is transformed back from the orthogonal modeling functions to the original physical quantities. Uncertainty bounds associated with each model parameter estimate are also computed. The multivariate orthogonal function modeling approach is applied to each of the six rigid-body degrees of freedom. [A more detailed discussion of the modeling module is given in \[17\], where it was first presented.](#)

Guidance: For the flights discussed here, the vehicle was overseen by either an R/C pilot or the real-time guidance algorithm [18]. An R/C safety pilot would take off and land the aircraft on each flight. While in control of the vehicle, the R/C pilot would fly the vehicle through various conditions, including stalls, and then engage the autonomous mode. In the R/C mode, the control law was not engaged and the pilot's commands are passed directly to the control surfaces. In the autonomous mode, the vehicle would fly back and forth between two waypoints following the guidance commands of flight path angle and heading.

Control: The primary purpose of the control law was to stabilize the vehicle and track the guidance commands. The control law had two mechanisms for handling disturbances and uncertainties within the system. In the short term, the controller would adapt to these unknowns like many traditional adaptive control architectures. However, in the long term, the controller would "learn" and adjust itself based on the real-time modeling results. Learning implies modifying the (desired) behavior or performance for the future based on retained knowledge from past experiences.

In order for the real-time modeling to learn the aerodynamics, a certain amount of perturbation is required, which is provided by the multi-sine inputs described above. However, the excitation provided by the multi-sine inputs is at odds with the controller's goal of stabilization. In simulation, we found that rejecting too much disturbance hindered the real-time aerodynamic modeling's ability to identify a model. During the L2F process, the conflicting goals of excitation for identification and stability for control must be balanced.

The controller is composed of a baseline controller and an adaptive augmentation. The baseline controller (a nonlinear dynamic inversion (NDI) controller) was designed based on the real-time learned model to maintain the natural frequency of the vehicle and provide sufficient damping for tracking the guidance commands. It uses the onboard aerodynamic model to effectively cancel out any undesirable dynamics. The adaptive augmentation assists the inner-loop NDI controller by comparing the measured response to the predicted response of the learned model and adapting the control signal to reduce any discrepancies. As mentioned before, periodic update of the learned model

naturally leads to a switched reference model for the adaptive control design. The next section addresses this adaptive control design problem.

III. \mathcal{L}_1 Adaptive Control for Uncertain Switched Linear Systems

In this section, we present the design and analysis of an \mathcal{L}_1 adaptive controller for a switched linear time-invariant (LTI) system subject to time-varying parametric uncertainties and disturbances. It is shown that, with fast adaptation, the closed-loop system with the \mathcal{L}_1 controller can behave arbitrarily close to a non-adaptive switched reference system that depends on perfect knowledge of the uncertainties. By starting with a stable reference system, the adaptive system is thus ensured to be stable. The work in this section differs from previous results [16] in that it adopts the piecewise-constant adaptive law [4, Section 3.2], which is more favorable for numerical implementation, as compared to the projection based adaptive law used in [16]. Additionally, we explicitly consider the potential (re-)initialization errors associated with the state predictor (e.g., due to measurement inaccuracies) and analyze the performance of the proposed adaptive control systems in the presence of such errors.

A. Problem Formulation

Consider the family of multi-input multi-output LTI systems whose state-space matrices are given by $\{(A_i, B_i, C_i) : i \in \mathcal{I}\}$, where \mathcal{I} denotes the index set. Let $\Sigma = \{\sigma : [0, \infty) \rightarrow \mathcal{I}\}$ denote the set of piecewise-constant switching signals. For a given switching signal $\sigma \in \Sigma$, define the following switched linear system subject to parametric uncertainties and disturbances:

$$\begin{aligned} \dot{x}(t) &= A_\sigma x(t) + B_\sigma (\omega_\sigma u(t) + \theta_\sigma^\top(t)x(t) + d_\sigma(t)), \\ y(t) &= C_\sigma x(t), \quad x(0) = x_0, \end{aligned} \tag{1}$$

where $x(t) \in \mathbb{R}^n$ is the system state, $u(t) \in \mathbb{R}^m$ is the system input, and $y(t) \in \mathbb{R}^m$ is the regulated system output, $A_\sigma \in \mathcal{A} \subset \mathbb{R}^{n \times n}$, $B_\sigma \in \mathcal{B} \subset \mathbb{R}^{n \times m}$, and $C_\sigma \in \mathcal{C} \subset \mathbb{R}^{m \times n}$ are the system matrices, $\omega_\sigma \in \Omega \subset \mathbb{R}^{m \times m}$, $\theta_\sigma(t) \in \Theta \subset \mathbb{R}^{n \times m}$, and $d_\sigma(t) \in \Delta \subset \mathbb{R}^m$ denote the unknown input gain, (time-varying) parametric uncertainties and external disturbances, respectively. Given a switching signal σ , we denote the sequence of finite switching time instant as $t_0, t_1, \dots, t_i, \dots$ with $t_0 = 0$. In the following, we explain the major assumptions.

Assumption 1. The sets \mathcal{A} , \mathcal{B} , \mathcal{C} , Ω , Θ , and Δ are compact, convex polytopes. We assume that each $B \in \mathcal{B} \subset \mathbb{R}^{n \times m}$ has full column rank $m (\leq n)$. Without loss of generality, the sets Θ and Δ are assumed to contain 0. Ω is assumed to be strictly diagonally dominant (which implies non-singularity), and, without loss of generality, is assumed to contain \mathbb{I}_m .

The control objective is to design a full-state feedback adaptive control law to ensure that $x(t)$ and $y(t)$ track the ideal signals $x_{\text{id}}(t)$ and $y_{\text{id}}(t)$ with quantifiable bounds on the transient and steady-state performances. The signals

$x_{\text{id}}(t)$ and $y_{\text{id}}(t)$ are given by the reference model that defines the desired closed-loop dynamics:

$$\begin{aligned} \dot{x}_{\text{id}}(t) &= A_{\sigma}x_{\text{id}}(t) + B_{\sigma}u_{\text{id}}(t), \quad x_{\text{id}}(0) = x_0, \\ u_{\text{id}}(t) &= k_{\sigma}r(t), \quad y_{\text{id}}(t) = C_{\sigma}x_{\text{id}}(t), \end{aligned} \tag{2}$$

where $r(t) \in \mathbb{R}^m$ is a given bounded piecewise-continuous reference signal, and $k_{\sigma} \in \mathbb{R}^{m \times m}$ is a time-varying feedforward gain for achieving a desired tracking performance. In the L2F framework, the ideal system will be determined based on the real-time modeling results and the NDI baseline controller.

Assumption 2. There exist symmetric matrices P_i ($i \in \mathcal{I}$) such that

$$\begin{aligned} P_i &\geq \mathbb{I}, \quad \forall i \in \mathcal{I}, \\ A_i^{\top}P_i + P_iA_i &\leq -\lambda P_i, \quad \forall i \in \mathcal{I}, \\ P_i &\leq \mu P_j, \quad \forall i, j \in \mathcal{I}, \end{aligned} \tag{3}$$

for some constants $\lambda > 0$ and $\mu \geq 1$. Moreover, the switching signal σ has a *dwell time*, $\tau_d \geq \frac{\ln(\mu)}{(1-a^*)\lambda}$, for arbitrary $a^* \in (0, 1)$. In other words, the switching times t_1, t_2, \dots satisfy the inequality $t_{k+1} - t_k \geq \tau_d$ for all $k = 0, 1, \dots$

Remark 1. Assumption 2 guarantees that the switched ideal dynamics (2) are stable. This can be easily proved following [19] using the switched Lyapunov function defined as $V_{\sigma}(x_{\text{id}}(t)) \triangleq x_{\text{id}}^{\top}(t)P_{\sigma}x_{\text{id}}(t)$. Essentially, the Lyapunov function is allowed to switch, at which time it may possibly grow in value by a fixed factor μ . By enforcing the dwell time condition based on the growth factor μ and the decay rate λ , the value of the Lyapunov will ultimately be decreasing over each switching interval. If we remove the dwell time constraint for the switching signal, i.e., allowing for arbitrary switching, a common Lyapunov function approach is usually employed to guarantee the stability. This can be seen as a special case of dwell-time switching with $\tau_d = 0$, $\mu = 1$. Note that the results derived in this section for dwell-time switching also hold for the more general case of average-dwell-time switching. For the details of (average) dwell-time switching, see [19].

B. \mathcal{L}_1 Adaptive Control Architecture

Figure 3 depicts the architecture of the proposed \mathcal{L}_1 controller and its interface with other modules of the L2F systems. The \mathcal{L}_1 controller is composed of three components, namely, a state predictor, an adaptive law and a low-pass filtered control law. For the subsequent introduction, let $B_{\sigma}^{\perp} \in \mathbb{R}^{n \times (n-m)}$ such that $B_{\sigma}^{\top}B_{\sigma}^{\perp} = 0$ and $\text{rank}(B_{\sigma}^{\vee}) = n$, where $B_{\sigma}^{\vee} \triangleq [B_{\sigma} \ B_{\sigma}^{\perp}]$.

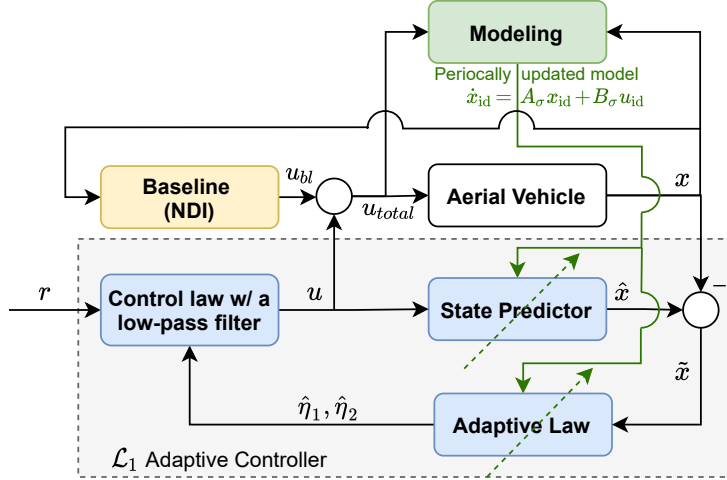


Fig. 3 Proposed \mathcal{L}_1 AC control architecture and its interface with the L2F system

State predictor: The state predictor is given by

$$\begin{aligned}\dot{\hat{x}}(t) &= A_\sigma \hat{x}(t) + B_\sigma (u(t) + \hat{\eta}_1(t)) + B_\sigma^\perp \hat{\eta}_2(t), \\ \hat{x}(t_i) &= \hat{x}_i,\end{aligned}\tag{4}$$

where \hat{x}_i is the reset value of the states, after switching to the sub-system i , and $\hat{\eta}_1(t)$ and $\hat{\eta}_2(t)$ are the estimated uncertainties.

Adaptive Law: The estimation of the (lumped) uncertainties $\hat{\eta}_1(t)$ and $\hat{\eta}_2(t)$ is updated by the following piecewise-constant adaptive law [4, Section 3.3]:

$$\begin{aligned}\begin{bmatrix} \hat{\eta}_1(t) \\ \hat{\eta}_2(t) \end{bmatrix} &= \begin{bmatrix} \hat{\eta}_1(jT_s) \\ \hat{\eta}_2(jT_s) \end{bmatrix}, \quad \forall t \in [jT_s, (j+1)T_s), \\ \begin{bmatrix} \hat{\eta}_1(jT_s) \\ \hat{\eta}_2(jT_s) \end{bmatrix} &= \begin{bmatrix} \mathbb{I}_m & 0 \\ 0 & \mathbb{I}_{n-m} \end{bmatrix} (B_\sigma^\vee)^{-1} A_\sigma \left(e^{-A_\sigma T_s} - \mathbb{I} \right)^{-1} \tilde{x}(jT_s),\end{aligned}\tag{5}$$

where T_s represents the estimation sampling time and $\tilde{x} \triangleq \hat{x} - x$ is the prediction error.

Control Law: The control input is computed through a feedback connection defined, in the Laplace domain, by

$$u(s) = -\frac{D_0(s)}{s} \mu(s),\tag{6}$$

where $\mu(s)$ is the Laplace transform of $\mu(t) \triangleq u(t) + \hat{\eta}_1(t) - k_\sigma r(t)$, and $D_0(s)$ is used to define a series of stable low-pass filters (defined later in (9)) that sufficiently guarantee the stability of a reference system introduced in Section III.C.

C. Stability and Performance analysis

In this section, we will analyze the stability and performance of the adaptive closed-loop system with the controller defined in (4)–(6), and derive performance bounds for both transient and steady-state phases. For this purpose, we first introduce a non-adaptive reference system that depends on the actual uncertainties (and thus is non-implementable), defined as

$$\begin{aligned} \dot{x}_{\text{ref}}(t) &= A_{\sigma}x_{\text{ref}}(t) + B_{\sigma}(\omega_{\sigma}u_{\text{ref}}(t) + \theta_{\sigma}^{\top}(t)x_{\text{ref}}(t) + d_{\sigma}(t)), \quad x_{\text{ref}}(0) = x_0, \\ u_{\text{ref}}(s) &= -\frac{D_0(s)}{s}\mu_{\text{ref}}(s), \end{aligned} \quad (7)$$

where $\mu_{\text{ref}}(s)$ is the Laplace transform of $\mu_{\text{ref}}(t) \triangleq \omega_{\sigma}u_{\text{ref}}(t) + \theta_{\sigma}^{\top}(t)x_{\text{ref}}(t) + d_{\sigma}(t) - k_{\sigma}r(t)$. *Such a reference system should not be conflated with the reference model used in other sections which characterizes the desired closed-loop dynamics (2).* The last equation in (7) is equivalent to*

$$u_{\text{ref}} = -\omega_{\sigma}^{-1}\mathcal{F}_{\sigma}(\xi_{\text{ref}} - r_k), \quad (8)$$

where $\xi_{\text{ref}}(t) \triangleq \theta_{\sigma}^{\top}(t)x_{\text{ref}}(t) + d_{\sigma}(t)$, $r_k(t) \triangleq k_{\sigma}r(t)$, and $\mathcal{F}_{\sigma} \triangleq \mathcal{F}_i$ when $\sigma(t) = i$. The (time-invariant) mapping \mathcal{F}_i has the transfer function form of

$$C_i(s) \triangleq \omega_i(s\mathbb{I} + D_0(s)\omega_i)^{-1}D_0(s) \quad (9)$$

which is a low-pass filter with DC gain equal to an identity matrix, i.e., $\mathcal{F}_i(0) = \mathbb{I}_m$. From (7) and (8), one can see that for each subsystem i , the reference control input tries to cancel the uncertainties within the bandwidth of the filter $\mathcal{F}_i(s)$.

Since the reference control law utilizes the true uncertainties, it is not implementable. The reference system is only introduced to help characterize the performance of the adaptive system by providing the ideal achievable performance. With the reference system, the performance of the adaptive system can be analyzed and quantified in four steps: (i) establish the stability of the reference system; (ii) quantify the difference between the state and input signals of the adaptive closed-loop system and those of the reference system; (iii) quantify the difference between the state and input signals of the reference system and those of the ideal system defined in (2); (iv) based on the results from (ii) and (iii), quantify the difference between the state and input signals of the adaptive system and those of the ideal system. Due to space limitations, we focus on the first two steps, while the last two steps are straightforward following existing results [4, 20]. To simplify the analysis, we also make the following reasonable assumption.

Assumption 3. The system switches take place at multiples of the sampling period, i.e., each switching time satisfies $t_i = kT_s$ for some k .

*We use the input-output mapping form instead of a transfer function form in (8) since the mapping is time-varying due to the existence of switching.

Remark 2. Within the L2F framework, the time instants for model switching are pre-determined, and therefore can be easily selected to satisfy the above assumption.

1. Stability of the Reference System

Letting (A_f, B_f, C_f, D_f) be a minimal realization of $D_0(s)$ with n_f states, the reference system dynamics in (7) can be rewritten in a state-space form as

$$\begin{aligned}\dot{\bar{x}}_{\text{ref}}(t) &= \bar{A}_\sigma \bar{x}_{\text{ref}}(t) + \bar{B}_\sigma d_\sigma(t) + \bar{E}_\sigma r(t), \quad \bar{x}(0) = \bar{x}_0, \\ u_{\text{ref}}(t) &= \bar{C} \bar{x}_{\text{ref}}(t),\end{aligned}\tag{10}$$

where $\bar{x}_{\text{ref}}(t) \triangleq [x_{\text{ref}}^\top(t), x_{f_1}^\top(t), x_{I_1}^\top(t)]^\top$, $\bar{x}_{\text{ref}}(0) \triangleq [x_0^\top, 0, 0]^\top$, $x_{f_1}(t) \in \mathbb{R}^{n_f}$ and $x_{I_1}(t) \in \mathbb{R}^m$ are the states of $D_0(s)$ and the integrator in (7), respectively, and

$$\bar{A}_\sigma \triangleq \left[\begin{array}{c|cc} A_\sigma + B_\sigma \theta_\sigma^\top & 0 & -B_\sigma \omega_\sigma \\ \hline B_f \theta_\sigma^\top & A_f & B_f \omega_\sigma \\ D_f \theta_\sigma^\top & C_f & D_f \omega_\sigma \end{array} \right], \quad \bar{B}_\sigma \triangleq \begin{bmatrix} B_\sigma \\ B_f \\ D_f \end{bmatrix}, \quad \bar{E}_\sigma \triangleq \begin{bmatrix} 0 \\ B_f k_\sigma \\ D_f k_\sigma \end{bmatrix}, \quad \bar{C} \triangleq \begin{bmatrix} 0 & 0 & -\mathbb{I} \end{bmatrix}.\tag{11}$$

Assumption 4. The parameters of $D_0(s)$ are designed such that there exist symmetric matrices $\bar{P}_i(\omega)$ ($i \in \mathcal{I}$) such that for all $(\theta, \omega) \in \Theta \times \Omega$ one has

$$\begin{aligned}\bar{P}_i &\geq \mathbb{I}, \quad \forall i \in \mathcal{I}, \\ \bar{A}_i^\top \bar{P}_i + \bar{P}_i \bar{A}_i &\leq -\lambda \bar{P}_i, \quad \forall i \in \mathcal{I}, \\ \bar{P}_i &\leq \mu \bar{P}_j, \quad \forall i, j \in \mathcal{I},\end{aligned}\tag{12}$$

for some constants $\lambda > 0$ and $\mu \geq 1$. Moreover, the switching signal satisfies the dwell time constraint

$$\tau_d \geq \frac{\ln(\mu)}{(1 - a^*)\lambda}\tag{13}$$

for arbitrary $a^* \in (0, 1)$.

Remark 3. Assumption 4 ensures that the control parameters of $D_0(s)$ are designed to stabilize the reference system (following [19] with the switched Lyapunov function $V_\sigma(\bar{x}_{\text{ref}}(t)) \triangleq \bar{x}_{\text{ref}}(t)^\top (t) \bar{P}_\sigma(\omega) \bar{x}_{\text{ref}}(t)$). As the gain of $D_0(s)$, corresponding to the bandwidth of the low-pass filter $\mathcal{F}_i(s)$ in (9), increases toward infinity, \mathcal{F}_i tends to become an identity matrix, \mathbb{I} , for all $i \in \mathcal{I}$. Thus, the effect of the uncertainty in the reference system vanishes. By substituting (8) into the dynamics of (7) with $\mathcal{F}_\sigma \equiv \mathbb{I}$, we see that the uncertainties are canceled perfectly, and the reference system becomes the ideal system, which is stable under Assumption 2. Therefore, by a perturbation argument, theoretically,

$D_0(s)$ can always be designed to stabilize the reference system if Assumption 2, which guarantees the stability of the ideal system, holds.

Note that when the reference system is stable, for a given reference signal r and initial condition x_0 , there exist constants ρ_r and ρ_{ur} such that

$$\begin{aligned}\|x_{\text{ref}}(t)\| &\leq \rho_r, \\ \|u_{\text{ref}}(t)\| &\leq \rho_{ur},\end{aligned}$$

for all $t \geq 0$ and for all admissible uncertainties. These bounds are independent of the adaptive law and state predictor.

2. Transient and Steady-State Performance

For the subsequent analysis, define $D_\theta \triangleq \max_{\theta \in \Theta} \|\theta\|$, $D_d \triangleq \max_{d \in \Delta} \|d\|$, $D_\omega \triangleq \max_{\omega \in \Omega} |\text{tr}(\omega - \mathbb{I})|$. Computing the prediction error dynamics from (1) and (4), we have

$$\begin{aligned}\dot{\tilde{x}}(t) &= A_\sigma \tilde{x}(t) + B_\sigma ((\mathbb{I} - \omega_\sigma)u + \hat{\eta}_1(t) - \eta_\sigma(t)) + B_\sigma^\perp \hat{\eta}_2(t), \\ \tilde{x}(t_i) &= \hat{x}(t_i) - x(t_i),\end{aligned}\tag{14}$$

where $x(t_i)$ is the state of the actual system when the i th switching happens at t_i , and

$$\eta_\sigma(t) = \theta_\sigma^\top(t)x(t) + d_\sigma(t).\tag{15}$$

For the following derivations, we first define some bounding functions:

$$\alpha_1(t) \triangleq \max_{i \in \mathcal{I}} \|e^{A_i t}\|, \quad \alpha_2(t) \triangleq \max_{i \in \mathcal{I}} \int_0^t \left\| e^{A_i(t-\tau)} A_i \left(e^{-A_i T_s} - \mathbb{I} \right)^{-1} \right\| d\tau, \quad \alpha_3(t) \triangleq \max_{i \in \mathcal{I}} \int_0^t \left\| e^{A_i(t-\tau)} B_i \right\| d\tau.$$

We further let

$$\bar{\alpha}_1(T_s) \triangleq \max_{t \in [0, T_s]} \alpha_1(t), \quad \bar{\alpha}_2(T_s) \triangleq \max_{t \in [0, T_s]} \alpha_2(t), \quad \bar{\alpha}_3(T_s) \triangleq \max_{t \in [0, T_s]} \alpha_3(t).$$

Notice that $\bar{\alpha}_1(T_s)$ and $\bar{\alpha}_2(T_s)$ are bounded. Moreover, $\bar{\alpha}_3(T_s)$ satisfies

$$\lim_{T_s \rightarrow 0} \bar{\alpha}_3(T_s) = 0.\tag{16}$$

Hence, for arbitrary $\delta_0 > 0$, T_s can be chosen such that

$$(\bar{\alpha}_1(T_s) + \bar{\alpha}_2(T_s) + 1) \bar{\alpha}_3(T_s) (D_\omega \rho_u + D_\theta \rho + D_\sigma) < \delta_0.\tag{17}$$

Assumption 5. We further assume that the re-initialization error at each switching instant satisfies

$$\|\tilde{x}(t_i)\| \leq \bar{\alpha}_3(T_s) (D_\omega \rho_u + D_\theta \rho + D_\sigma). \quad (18)$$

Remark 4. Note that the states of the actual system are measurable; therefore, theoretically, the re-initialization error could be made zero since the states of the state predictor can be set to an arbitrary value. In practice, the states of the actual system may not be measured exactly, e.g., due to measurement error/noise. We thus consider non-zero re-initialization error to account for the practical scenario.

The following lemma establishes a bound on the prediction error, which can be arbitrarily small under some conditions on the estimation sampling time T_s and re-initialization error.

Lemma 1. Assume $\|x(t)\| \leq \rho$ and $\|u(t)\| \leq \rho_u$ for all $0 \leq t \leq \tau$ for some positive constants ρ and ρ_u . Given an arbitrary positive constant δ_0 , if T_s is chosen to satisfy (17) and Assumption 5 holds, then $\|\tilde{x}(t)\| < \delta_0$ for all $0 \leq t \leq \tau$.

Proof. For arbitrary $j \geq 0$, it follows from the prediction error dynamics in (14) that for the i th system on the interval $t \in [0, T_s]$ with $t_i + jT_s + t \leq \min(t_{i+1}, \tau)$, we have

$$\tilde{x}(t_i + jT_s + t) = \underbrace{e^{A_i t} \tilde{x}(t_i + jT_s)}_{\zeta_1(t_i + jT_s + t)} + \underbrace{\int_0^t e^{A_i(t-\tau)} B_i^\vee \begin{bmatrix} \hat{\eta}_1(\tau) \\ \hat{\eta}_2(\tau) \end{bmatrix} d\tau}_{\zeta_2(t_i + jT_s + t)} + \underbrace{\int_0^t e^{A_i(t-\tau)} B_i ((\mathbb{I} - \omega_i)u_i - \eta_i(\tau)) d\tau}_{\zeta_2(t_i + jT_s + t)}. \quad (19)$$

Substituting the adaptive law in (5) and evaluating at $t = T_s$, we get

$$\begin{aligned} \zeta_1(t_i + jT_s + T_s) &= e^{A_i T_s} \tilde{x}(jT_s) + \int_0^{T_s} e^{A_i(T_s-\tau)} B_i^\vee (B_i^\vee)^{-1} A_i (e^{-A_i T_s} - \mathbb{I})^{-1} \tilde{x}(jT_s) d\tau = 0 \\ \zeta_2(t_i + jT_s + T_s) &= \int_0^{T_s} e^{A_i(T_s-\tau)} B_i ((\mathbb{I} - \omega_i)u - \eta_i(\tau)) d\tau \leq \bar{\alpha}_3(T_s) (D_\omega \rho_u + D_\theta \rho + D_\sigma) \end{aligned}$$

Thus, for arbitrary $j \geq 0$, it follows that $\tilde{x}(t_i + jT_s) \leq \bar{\alpha}_3(T_s) (D_\omega \rho_u + D_\theta \rho + D_\sigma)$. From (19) and substituting the adaptive laws (5), we obtain

$$\begin{aligned} \|\tilde{x}(t_i + jT_s + t)\| &\leq \|e^{A_i t} \tilde{x}(t_i + jT_s)\| + \left\| \int_0^t e^{A_i(t-\tau)} A_i (\mathbb{I} - e^{-A_i T_s})^{-1} \tilde{x}(t_i + jT_s) d\tau \right\| + \left\| \int_0^t -e^{A_i(t-\tau)} B_i d_\sigma(\tau) d\tau \right\| \\ &\leq \alpha_1(t) \|\tilde{x}(jT_s)\| + \alpha_2(t) \|\tilde{x}(jT_s)\| + \alpha_3(t) (D_\omega \rho_u + D_\theta \rho + D_d) \\ &\leq (\bar{\alpha}_1(T_s) + \bar{\alpha}_2(T_s) + 1) \bar{\alpha}_3(T_s) (D_\omega \rho_u + D_\theta \rho + D_d) < \delta_0, \end{aligned}$$

where the last inequality is due to (17). This bound holds uniformly for all $i \in \mathcal{I}$, for all $j \geq 0$, and for all $t \in [0, T_s]$.

Therefore, $\|\tilde{x}(t)\| < \delta_0$ for all $0 \leq t \leq \tau$. The proof is complete.

Remark 5. If the re-initialization is not done such that Assumption 5 is met, we see from the proof of Lemma 1 that the effect of this re-initialization error disappears after one time step. A similar observation is made in \mathcal{L}_1 adaptive control of linear parameter-varying systems [20] and adaptive estimation for nonlinear systems [21]. Moreover, by allowing the state predictor to evolve continuously, i.e., no re-initialization, Assumption 5 will automatically be satisfied all the time except perhaps at the initial time $t = 0$.

From the prediction error dynamics in (14), we have that for all $i \in \mathcal{I}$,

$$\tilde{\eta}_i(t) = B_i^\dagger (\dot{\tilde{x}}(t) - A_i \tilde{x}(t) - B_i^\perp \hat{\eta}_2(t)) = B_i^\dagger \dot{\tilde{x}}(t) - B_i^\dagger A_i \tilde{x}(t), \quad (20)$$

where B_i^\dagger is the pseudo-inverse of B_i . Letting $y_{f_2}(t)$ be the output after passing $B_i^\dagger \dot{\tilde{x}}(t)$ through the filter $\mathcal{F}_i(s)$ (defined in (9)), we get

$$\begin{bmatrix} \dot{x}_{f_2}(t) \\ \dot{x}_{I_2}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} A_f & B_f \omega_i \\ C_f & D_f \omega_i \end{bmatrix}}_{\triangleq \bar{F}_i} \begin{bmatrix} x_{f_2}(t) \\ x_{I_2}(t) \end{bmatrix} + \underbrace{\begin{bmatrix} B_f \\ D_f \end{bmatrix}}_{\triangleq \bar{B}_f} B_i^\dagger \dot{\tilde{x}}(t), \quad y_{f_2}(t) = \underbrace{\begin{bmatrix} C_f & D_f \omega_i \end{bmatrix}}_{\triangleq \bar{L}_i} \begin{bmatrix} x_{f_2}(t) \\ x_{I_2}(t) \end{bmatrix} + D_f B_i^\dagger \dot{\tilde{x}}(t) - \gamma_i(t), \quad (21)$$

where $x_{f_2}(t) \in \mathbb{R}^{n_f}$ and $x_{I_2}(t) \in \mathbb{R}^m$ are the states of $D_0(s)$ and the integrator in (9), respectively, and $\gamma_i(t) = \bar{C}_f e^{\bar{F}_i(t-t_i)} \bar{B}_f B_i^\dagger \tilde{x}(t_i)$ with $\bar{C}_f = [0 \quad \mathbb{I}]$. We note that $\gamma_i(t)$ can be bounded by

$$\|\gamma_i(t)\| \leq \kappa_\gamma e^{\Lambda_{\bar{F}}(t-t_i)} \|\tilde{x}(t_i)\|, \quad (22)$$

where $\kappa_\gamma = \max_{i \in \mathcal{I}} \|\bar{C}_f \bar{B}_f B_i^\dagger\|$ and $\Lambda_{\bar{F}} = \max_{i \in \mathcal{I}} \lambda_{\max}(\bar{F}_i)$, with $\lambda_{\max}(\cdot)$ representing the largest eigenvalue.

We next characterize the difference between the responses of the adaptive closed-loop system and of the reference system. Note that u in (6) can be equivalently written as

$$u = -\omega_\sigma^{-1} \mathcal{F}_\sigma (\eta_\sigma + \tilde{\eta}_\sigma - k_\sigma r), \quad (23)$$

where $\eta_\sigma(t)$ is defined in (15), \mathcal{F}_σ is introduced in (8) and $\tilde{\eta}_\sigma(t) \triangleq \hat{\eta}_1(t) - \eta_\sigma(t) - (\omega_\sigma - \mathbb{I})u(t)$. Applying the control law (23) to the system dynamics (1), we can rewrite the adaptive closed-loop system similarly to (10) as

$$\begin{aligned} \dot{\tilde{x}}(t) &= \bar{A}_\sigma \tilde{x}(t) + \bar{B}_\sigma d_\sigma(t) + \bar{E}_\sigma r(t) + \begin{bmatrix} 0 & B_f^\top & D_f^\top \end{bmatrix}^\top \tilde{\eta}_\sigma(t), \\ u(t) &= \bar{C} \tilde{x}(t), \end{aligned} \quad (24)$$

where $\tilde{x}(t) \triangleq [x^\top(t), x_{f_1}^\top(t), x_{I_1}^\top(t)]^\top$, $\tilde{x}^\top(0) \triangleq [x_0^\top, 0, 0]^\top$, x_{f_1} and x_{I_1} have the same meaning as in (10), \bar{A}_σ , \bar{B}_σ , \bar{C} and \bar{E}_σ are defined in (11).

The difference between the state and input of the actual system and those of the reference system is denoted as

$e(t) \triangleq x_{\text{ref}}(t) - x(t)$ and $e_u(t) \triangleq u_{\text{ref}}(t) - u(t)$. With the definition of

$$\bar{e}(t) \triangleq [e^\top(t), x_{f_1}^\top(t), x_{I_1}^\top(t)]^\top, \quad (25)$$

the equations (10) and (24) imply that

$$\begin{aligned} \dot{\bar{e}}(t) &= \bar{A}_\sigma \bar{e}(t) - \begin{bmatrix} 0 & B_f^\top & D_f^\top \end{bmatrix}^\top \bar{\eta}_\sigma(t), \quad \bar{e}(0) = 0, \\ e_u(t) &= \bar{C} \bar{e}(t) \end{aligned} \quad (26)$$

Applying (20) and (21) to the error dynamics in (26) over the time interval $t \in [t_i, t_{i+1})$, we have

$$\begin{aligned} \begin{bmatrix} \dot{\bar{e}}(t) \\ \dot{\bar{x}}_{f_2}(t) \end{bmatrix} &= \begin{bmatrix} \bar{A}_i & \bar{H}_i \\ 0 & \bar{F}_i \end{bmatrix} \begin{bmatrix} \bar{e}(t) \\ \bar{x}_{f_2}(t) \end{bmatrix} + \begin{bmatrix} \bar{J}_i \\ \bar{G}_i \end{bmatrix} \bar{x}(t) + \begin{bmatrix} \bar{B}_i \\ 0 \end{bmatrix} \gamma_i(t), \\ \begin{bmatrix} e_u(t) \end{bmatrix} &= \begin{bmatrix} \bar{C} & \bar{L}_i \end{bmatrix} \begin{bmatrix} \bar{e}(t) \\ \bar{x}_{f_2}(t) \end{bmatrix} + \begin{bmatrix} -D_f B_i^\dagger \end{bmatrix} \bar{x}(t), \end{aligned} \quad (27)$$

where $\bar{x}_{f_2}(t) = [x_{f_2}^\top(t), x_{I_2}^\top(t)]^\top$, $\bar{e}(t) = [e^\top(t), \bar{x}_{f_1}^\top(t)]^\top$, and

$$\bar{H}_i = \begin{bmatrix} -B_i C_f & -B_i D_f \omega_i \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \bar{J}_i = \begin{bmatrix} -D_f B_i^\dagger \\ -B_f B_i^\dagger A_i \\ -D_f B_i^\dagger A_i \end{bmatrix}, \quad \bar{B}_i = \begin{bmatrix} B_i \omega_i \\ 0 \\ 0 \end{bmatrix}, \quad \bar{G}_i = -\bar{B}_f B_i^\dagger A_i.$$

We partition \bar{P}_i in (12) along the same partition as \bar{A}_i in (11), which leads to

$$\bar{P}_i = \left[\begin{array}{c|c} P_i & R_i \\ \hline R_i^\top & S_i \end{array} \right].$$

Further defining $Q_i \triangleq S_i - R_i^\top P_i^{-1} R_i$, we obtain

$$\begin{aligned} Q_i &\geq \mathbb{I}, \quad \forall i \in \mathcal{I}, \\ \bar{F}_i^\top Q_i + Q_i \bar{F}_i &\leq -\lambda Q_i, \quad \forall i \in \mathcal{I}, \\ Q_i &\leq \mu Q_j, \quad \forall i, j \in \mathcal{I}. \end{aligned} \quad (28)$$

Let a scalar $\nu > 0$ satisfy

$$-\lambda a \bar{P}_i + \frac{1}{\nu \lambda a} \bar{P}_i \bar{H}_i Q_i^{-1} \bar{H}_i^\top \bar{P}_i < 0 \quad \forall i \in \mathcal{I}, \quad (29)$$

for some $a \in (0, a^*)$, where a^* was introduced in (13). Note that such a ν always exists since $\bar{P}_i > 0$.

Finally, let

$$\delta_1 \triangleq \sqrt{\left(\mu \left(1 - \mu^{\frac{a-a^*}{1-a^*}}\right)^{-1} + 1\right) \frac{g}{(1-a)\lambda} (1 + \kappa_y^2) \delta_0}, \quad \delta_2 \triangleq \max_{i \in \mathcal{I}} \left\| \left[\bar{C} \quad \bar{L}_i / \sqrt{\nu} \right] \right\| \delta_1 + \max_{i \in \mathcal{I}} \left\| D_f B_i^\dagger \right\| \delta_0. \quad (30)$$

We are now ready to state the following theorem which quantifies the performance of the adaptive control system, both transient and steady-state, in terms of its states and inputs, as compared to those of the reference system.

Theorem 1. *Consider the closed-loop adaptive system with the \mathcal{L}_1 controller defined via (4)–(6) and the closed-loop reference system (7). Suppose that there exist $D_0(s)$ and $\bar{P}_i(\omega)$ ($i \in \mathcal{I}$) and some constants $\lambda > 0$ and $\mu \geq 1$ such that the inequalities in (12) hold for all $(\theta, \omega) \in \Theta \times \Omega$; furthermore, suppose that the dwell time satisfies (13) and the sampling time T_s satisfies (17). Then, for all $t \geq 0$, we have*

$$\|\bar{x}(t)\| < \delta_0, \quad (31)$$

$$\|x(t)\| \leq \rho \triangleq \rho_r + \delta_1, \quad (32)$$

$$\|u(t)\| \leq \rho_u \triangleq \rho_{ur} + \delta_2, \quad (33)$$

$$\|x_{\text{ref}}(t) - x(t)\| \leq \delta_1, \quad (34)$$

$$\|u_{\text{ref}}(t) - u(t)\| \leq \delta_2, \quad (35)$$

where δ_0 is introduced in (17) and δ_i ($i = 1, 2$) are defined in (30).

Proof. Suppose that one of the bounds (34) or (35) does not hold. Since $e(t)$ and $e_u(t)$ are continuous with $e(0) = 0$ and $e_u(0) = 0$, there exists a τ such that

$$\|e(\tau)\| = \delta_1 \text{ or} \quad (36)$$

$$\|e_u(\tau)\| = \delta_2, \quad (37)$$

and

$$\|e(t)\| < \delta_1, \quad \|e_u(t)\| < \delta_2, \quad \forall t \in [0, \tau).$$

It follows then that for all $t \in [0, \tau)$, one has

$$\|x(t)\| \leq \rho \triangleq \rho_r + \delta_1,$$

$$\|u(t)\| \leq \rho_u \triangleq \rho_{ur} + \delta_2.$$

Applying Lemma 1, for all $t \in [0, \tau)$, we have

$$\|\tilde{x}(t)\| < \delta_0. \quad (38)$$

Let $V_i(t) = \bar{e}^\top(t) \bar{P}_i \bar{e}(t) + \nu \bar{x}_{f_2}^\top(t) Q_i \bar{x}_{f_2}(t)$ on the time interval $t \in [t_i, t_{i+1})$, where \bar{e} is defined in (25), and ν comes from (29). Differentiating $V_i(t)$ along the system trajectories, we have

$$\begin{aligned} \dot{V}_i(t) &= \nu \bar{x}_{f_2}^\top(t) (\bar{F}_i^\top Q_i + Q_i \bar{F}_i) \bar{x}_{f_2}(t) + 2\nu \bar{x}_{f_2}^\top(t) Q_i \bar{G}_i \tilde{x}(t) \\ &\quad + \bar{e}^\top(t) (\bar{A}_i^\top \bar{P}_i + \bar{P}_i \bar{A}_i) \bar{e}(t) + 2\bar{e}^\top(t) \bar{P}_i \bar{H}_i \bar{x}_{f_2}(t) + 2\bar{e}^\top(t) \bar{P}_i \bar{J}_i \tilde{x}(t) + 2\bar{e}^\top(t) \bar{P}_i \bar{B}_i \gamma_i(t) \\ &\leq \begin{bmatrix} \bar{e}^\top(t) & \bar{x}_{f_2}^\top(t) & \tilde{x}^\top(t) & \gamma_i^\top(t) \end{bmatrix} \begin{bmatrix} -\lambda \bar{P}_i & \bar{P}_i \bar{H}_i & \bar{P}_i \bar{J}_i & \bar{P}_i \bar{B}_i \\ \bar{H}_i^\top \bar{P}_i & -\nu \lambda Q_i & \nu Q_i \bar{G}_i & 0 \\ \bar{J}_i^\top \bar{P}_i & \nu \bar{G}_i^\top Q_i & 0 & 0 \\ \bar{B}_i^\top \bar{P}_i & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} \bar{e}(t) \\ \bar{x}_{f_2}(t) \\ \tilde{x}(t) \\ \gamma_i(t) \end{bmatrix} \\ &\leq -(1-a)\lambda V_i(t) + g \|\tilde{x}(t)\|^2 + g \|\gamma_i(t)\|^2, \end{aligned} \quad (39)$$

where the last line follows from square completion, and the scalar g is given by

$$g \triangleq \left\| \begin{bmatrix} \bar{J}_i^\top \bar{P}_i & \nu \bar{G}_i^\top Q_i \\ \bar{B}_i^\top \bar{P}_i & 0 \end{bmatrix} \begin{bmatrix} -\lambda a \bar{P}_i & \bar{P}_i \bar{H}_i \\ \bar{H}_i^\top \bar{P}_i & -\nu \lambda a Q_i \end{bmatrix}^{-1} \begin{bmatrix} \bar{P}_i \bar{J}_i & \bar{P}_i \bar{B}_i \\ \nu Q_i \bar{G}_i & 0 \end{bmatrix} \right\|.$$

Integrating the last inequality in (39), we have:

$$V_i(t) \leq V_i(t_i) e^{-(1-a)\lambda(t-t_i)} + \int_{t_i}^t e^{-(1-a)\lambda(t-\tau)} g \left(\|\tilde{x}(t)\|^2 + \|\gamma_i(t)\|^2 \right) d\tau.$$

Substituting the bound for $\|\gamma_i(t)\|$ from (22) and the bound for $\|\tilde{x}(t)\|$ from (38), we have

$$\begin{aligned} V_i(t) &< V_i(t_i) e^{-(1-a)\lambda(t-t_i)} + \int_{t_i}^t e^{-(1-a)\lambda(t-\tau)} g \left(1 + \kappa_\gamma^2 e^{2\Lambda_{\bar{F}}(t-t_i)} \right) \delta_0^2 d\tau \\ &< \mu V_{i-1}(t_i) e^{-(1-a)\lambda(t-t_i)} + \frac{g}{(1-a)\lambda} \left(1 + \kappa_\gamma^2 \right) \delta_0^2. \end{aligned} \quad (40)$$

If at any switching time t_i

$$V_{i-1}(t_i) < \left(1 - \mu \frac{a-a^*}{1-a^*} \right)^{-1} \frac{g}{(1-a)\lambda} \left(1 + \kappa_\gamma^2 \right) \delta_0^2,$$

then it follows from (40) and the dwell time constraint (12) that

$$\begin{aligned} V_i(t_{i+1}) &< \left(1 + \mu \left(1 - \mu^{\frac{a-a^*}{1-a^*}}\right)^{-1} \mu^{-\frac{1-a}{1-a^*}}\right) \frac{g}{(1-a)\lambda} (1 + \kappa_\gamma^2) \delta_0^2 \\ &< \left(1 - \mu^{\frac{a-a^*}{1-a^*}}\right)^{-1} \frac{g}{(1-a)\lambda} (1 + \kappa_\gamma^2) \delta_0^2, \end{aligned}$$

and it holds for all t_j with $j \geq i$. Since $V_i(0) = 0$, it holds for all switching times. Therefore, from (40) we have

$$\begin{aligned} \left\| \begin{bmatrix} \bar{e}(t) \\ \sqrt{v} \bar{x}_{\mathcal{F}}(t) \end{bmatrix} \right\|^2 &\leq V_i(t) < \left(\mu \left(1 - \mu^{\frac{a-a^*}{1-a^*}}\right)^{-1} e^{-(1-a)\lambda(t-t_i)} + 1 \right) \frac{g}{(1-a)\lambda} (1 + \kappa_\gamma^2) \delta_0^2 \\ &< \left(\mu \left(1 - \mu^{\frac{a-a^*}{1-a^*}}\right)^{-1} + 1 \right) \frac{g}{(1-a)\lambda} (1 + \kappa_\gamma^2) \delta_0^2 = \delta_1^2, \end{aligned}$$

and from the definition of e_u in (26), we have

$$\|e_u(t)\| < \max_{i \in \mathcal{I}} \left\| \begin{bmatrix} \bar{C} & \bar{L}_i / \sqrt{v} \end{bmatrix} \right\| \delta_1 + \max_{i \in \mathcal{I}} \|D_f B_i^\dagger\| \delta_0 = \delta_2.$$

This contradicts (36), and thus (34) and (35) hold. Applying the triangle inequality then gives (32) and (33). Finally, (31) follows from Lemma 1. The proof is complete.

To summarize, if the reference system is stable and the stated assumptions are met, Theorem 1 guarantees that the adaptive system follows the reference system within some a priori bounds. If the re-initialization errors do not satisfy Assumption 3, then, as stated in Remark 5, the effect of the re-initialization error only lasts for one time step. The size of the bounds in Theorem 1 are dependent on the choice of the time step T_s . By decreasing T_s , under the adaptive control law in (4), (5), (6), the system can be made arbitrarily close to the stable reference system.

IV. Controller Design for Learn-to-Fly

In this section, we introduce the details of the controller design for the L2F system, including how the switched \mathcal{L}_1 controller presented in Section III is incorporated. Before proceeding, we present a table listing the nomenclature used in this section. Any reuse of symbols in this section should not be confused with those used in the derivation and analysis of the \mathcal{L}_1 control architecture in Section III.

Table 1 Aircraft Nomenclature

Symbol	Meaning	Symbol	Meaning
V	Total Airspeed	L	Lift Force
α	Angle of Attack	Y	Side Force
β	Angle of Sideslip	ϕ	Roll Angle
p	Roll Rate	θ	Pitch Angle
q	Pitch Rate	χ	Ground Track Angle
r	Yaw Rate	γ	Flight Path Angle
ω	Angular Rate Vector $[p \ q \ r]^\top$	μ	Bank Angle
M	External Aerodynamic Moment Vector	m	Vehicle Mass
M_δ	Moment Produced by Control Surfaces	I	Inertia Matrix
g	Acceleration Due to Gravity		

A. Mathematical Model

To begin, we define the mathematical model used to develop the control law. We start with the standard equations of motion for a rigid-body aircraft [22] with no thrust:

$$\begin{aligned}
 \dot{\chi} &= \frac{1}{mV \cos(\gamma)} (L \sin(\mu) + Y \cos(\mu) \cos(\beta)), \\
 \dot{\gamma} &= \frac{1}{mV} (L \cos(\mu) - Y \sin(\mu) \cos(\beta) - mg \cos(\gamma)), \\
 \dot{\phi} &= p + \tan(\theta) (q \sin(\phi) + r \cos(\phi)), \\
 \dot{\beta} &= -\cos(\alpha)r + \sin(\alpha)p + \frac{1}{mV} (Y \cos(\beta) + mg \cos(\gamma) \sin(\mu)), \\
 \dot{\alpha} &= q - \tan(\beta) (\cos(\alpha)p + \sin(\alpha)r) + \frac{1}{mV \cos(\beta)} (-L + mg \cos(\gamma) \cos(\mu)), \\
 \dot{\theta} &= q \cos(\phi) - r \sin(\phi), \\
 \dot{\omega} &= I^{-1} (M + M_\delta) - I^{-1} (\omega \times I \omega).
 \end{aligned} \tag{41}$$

If we assume small angles (α, β, γ) and small rates (p, r) with negligible side force ($Y = 0$) and steady 1-g flight

($L = mg/\cos(\phi)$), the dynamics in (41) can be simplified as

$$\begin{aligned}
\dot{\chi} &= \frac{g}{V} \tan(\phi), \\
\dot{\gamma} &= 0, \\
\dot{\phi} &= p + \tan(\theta) (q \sin(\phi) + r \cos(\phi)), \\
\dot{\beta} &= -r + \frac{g}{V} \sin(\phi), \\
\dot{\alpha} &= q - \frac{g}{V} \sin(\phi) \tan(\phi), \\
\dot{\theta} &= q \cos(\phi) - r \sin(\phi), \\
\dot{\omega} &= I^{-1} (M + M_{\delta}) - I^{-1} (\omega \times I \omega).
\end{aligned} \tag{42}$$

This simplified model was used to develop the control laws. Originally, the baseline control law was designed for an unpowered glider-type aircraft. In addition, during these tests, the real-time modeling software was not able to separate the thrust effects from the aerodynamic effects. For these reasons, despite the test vehicle having thrust, the control law does not consider its effects. The velocity was loosely controlled by a simple proportional-integral law feeding the velocity error back to the throttle. The feedback loop was unknown to the controller and can be thought of as unmodeled dynamics. As an example, if we look at the angle of attack dynamics in the presence of thrust T , we can see how the interaction might skew the aerodynamic model:

$$\dot{\alpha} = q - \tan(\beta) (\cos(\alpha)p + \sin(\alpha)r) + \frac{1}{mV \cos(\beta)} (-L + mg \cos(\gamma) \cos(\mu) - T_x \sin(\alpha) + T_z \cos(\alpha)).$$

Assuming the thrust is axially aligned (i.e., $T_z = 0$), the modeled lift force will contain both the actual lift force L and a thrust component $T_x \sin(\alpha)$. With the thrust loop closed based on the velocity, the modeled lift force now contains an artificial dependence on speed, angle of attack, and the dynamics of motor thrust.

Since this initial flight testing, the capabilities of the real-time aerodynamic modeling have expanded to capture thrust effects by including modeling terms based on the advance ratio (similar to the use of corrected engine rotational speed in [23]). Unfortunately, further flight tests examining the corresponding effect on the performance of the baseline and adaptive controllers have not been performed.

B. Baseline Controller Design

As mentioned in Section II, the baseline controller is a nonlinear dynamic inversion (NDI) controller. To apply NDI to the dynamics in (42), we sequentially generate commands based on the time-scale separation of the variables. The first loop converts the guidance commands of ground track angle χ and flight path angle γ into attitude commands for the roll angle ϕ and pitch angle θ , respectively. The desired derivatives are chosen to be proportional to the error

between the variables and its command in order to produce a linear tracking response. Inverting the χ dynamics from (42) in this manner produces a roll command given by

$$\phi_{cmd} = \text{Sat}_{-45^\circ}^{+45^\circ} \left(\arctan \left(\frac{V}{g} K_\chi (\chi_{cmd} - \chi) \right) \right),$$

where we have additionally limited the roll command between $\pm 45^\circ$. Since the flight path angle dynamics in (42) are constant, we use a kinematic relationship between flight path angle γ and pitch angle θ from the z -component of the velocity vector to directly solve for a commanded pitch angle:

$$\sin(\gamma_{cmd}) = a_1 \cos(\theta_{cmd}) + a_2 \sin(\theta_{cmd}),$$

where $a_1 = \cos(\alpha) \cos(\beta)$ and $a_2 = \sin(\phi) \sin(\beta) + \cos(\phi) \sin(\alpha) \cos(\beta)$ [24].

The second loop converts the angle commands into angular rate commands, with ϕ , θ , and β being mapped to commands for p , q , and r , respectively. Once again seeking a linear tracking response, from (42), the commands are

$$\begin{aligned} p_{cmd} &= K_\phi (\phi_{cmd} - \phi) - \tan(\theta) (q \sin(\phi) + r \cos(\phi)), \\ q_{cmd} &= \frac{1}{\cos(\phi)} (K_\theta (\theta_{cmd} - \theta) + r \sin(\phi)), \\ r_{cmd} &= -K_\beta (\beta_{cmd} - \beta) - \frac{g}{V} \sin(\phi). \end{aligned}$$

The final loop converts these angular rate commands into moment commands:

$$M_{\delta,cmd} = IK_\omega (\omega_{cmd} - \omega) - (\hat{M} - \omega \times I\omega),$$

where $K_\omega = \text{diag}([K_p, K_q, K_r])$, $\omega_{cmd} = [p_{cmd}, q_{cmd}, r_{cmd}]^\top$, and an estimate of the current moment on the aircraft \hat{M} , as determined by the real-time aerodynamic modeling discussed in Section II and detailed in [17], is used in place of a measurement.

C. Model and Inflight Control Law Update from Online Learning

The control parameters K_θ , K_q , K_β , K_r , K_ϕ , and K_p were selected, based on the updated model, to maintain the natural frequency of the vehicle while supplying sufficient damping. Assuming the dynamic inversion is capable of removing the nonlinear effects, which would leave only the desired linear dynamics, the longitudinal dynamics would look like

$$\begin{bmatrix} \dot{\theta} \\ \dot{q} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -K_q K_\theta & -K_q \end{bmatrix} \begin{bmatrix} \theta \\ q \end{bmatrix} + \begin{bmatrix} 0 \\ K_q K_\theta \end{bmatrix} \theta_{cmd},$$

Table 2 Desired frequency calculations

Degree of freedom	Equation
Roll	$\omega_n = \sqrt{\left \frac{\bar{q}Sb}{2I_{xx}} C_{l_{\delta a}} \right }$
Pitch	$\omega_n = \sqrt{\left \frac{\bar{q}S\bar{c}}{I_{yy}} C_{m_{\alpha}} \right }$
Yaw	$\omega_n = \sqrt{\left \frac{\bar{q}Sb}{I_{zz}} C_{n_{\beta}} \right }$

which is a canonical description of a second-order system. The control parameters can then be written in terms of the natural frequency ω_n and damping ζ of the second-order system

$$K_q = 2\zeta\omega_n$$

$$K_{\theta} = \omega_n/2\zeta.$$

A similar description can be given for the roll and yaw dynamics.

The desired natural frequency was determined based on the linearized real-time modeling results, as described in Table 2, and the desired damping ratio was selected to be 0.8. The calculations also contain a number of physical vehicle parameters that are known a priori, such as the wing reference area S , the wing span b , the wing mean aerodynamic chord \bar{c} , and the principle components of the inertia matrix I_{xx} , I_{yy} , and I_{zz} . The non-linear, global non-dimensional model was updated at a rate of 5 Hz by the real-time modeling. The linearization of the non-dimensional model was then computed based on the desired operating point. From Table 2, we see that the desired natural frequencies (desired dynamics) were parameterized by the dynamic pressure \bar{q} during dimensionalization. Given the 5 Hz update rate, we reasonably assumed that the dynamic pressure remains constant over that interval; thus, we consider the system to be comprised of switched linear systems as discussed in Section III. For comparison, the flight control law implementation ran at a rate of 50 Hz.

D. \mathcal{L}_1 Adaptive Augmentation

The adaptive portion of the controller is used to assist the control loop regulating the angular rates. The desired dynamics are defined as the dynamic inversion component successfully canceling the nonlinearities and enforcing the ideal linear response:

$$\dot{\omega}(t) = -K_{\omega}\omega(t) + K_{\omega}\omega_{cmd}(t).$$

This can be represented in the form of (2):

$$\begin{aligned}\dot{x}_{id}(t) &= A_{\sigma}x_{id}(t) + B_{\sigma}u_{id}(t), & x_{id}(0) &= x_0, \\ u_{id}(t) &= k_{\sigma}r(t), & y_{id}(t) &= C_{\sigma}x_{id}(t),\end{aligned}$$

where we now transition back into the adaptive control nomenclature with the state x representing the angular velocity ω , and the reference r representing the commanded angular velocity ω_{cmd} . The system matrices A_{σ} and B_{σ} are given by $-K_{\omega}$ and \mathbb{I} , respectively, and k_{σ} and C_{σ} are equal to K_{ω} and \mathbb{I} , respectively. The values for K_{ω} are updated (switched) based on the real-time modeling results, as discussed in the previous subsection.

Due to the simplifying assumptions, model uncertainties, external disturbances, etc., the dynamic inversion controller will not be able to exactly achieve its goal. We thus arrive at the uncertain dynamics considered in (1) and apply the adaptive controller discussed in Section III.

V. Flight Test Results

The flight tests described here were performed by NASA Langley Research Center on the E1 flight test vehicle, shown in Figure 4. The vehicle is able to be flown in three modes: safety mode, pilot command pass-through mode, and full autonomy mode. In the safety mode, the vehicle is remotely piloted in a stick-to-surface fashion. This allows the R/C pilot to takeoff and land the vehicle, as well as to intervene at any time in order to preserve the safety of the vehicle, test crew, or airspace. The pilot command pass-through mode allows signals to be injected on top of the R/C pilot commands. In the nominal flight configurations, this allows the introduction of the PTIs during piloted operation. For the destabilized configurations, it is also used to introduce a destabilizing feedback loop to redundant control surfaces. For more details on the destabilization and the modes in general, see [25]. Finally, the full autonomy mode engages the onboard guidance, modeling, and control algorithms with the injected PTIs and any destabilizing feedback inputs. The pilot has no control of the vehicle in this mode.



Fig. 4 E1 Flight Test Vehicle

A brief summary of what was varied between flights is included in Table 3, along with any corresponding figure numbers. A more detailed description of the flights and discussion of the flight data is provided in the next two

subsections.

Table 3 Flight Test Summary

Configuration	Controller	Learning	Figures
Nominal Dynamics	NDI Only	Pilot Mode	Fig. 5
Nominal Dynamics	NDI+ \mathcal{L}_1	Pilot Mode	Fig. 6, 7
Pitch Destabilized (-10% static margin)	NDI+ \mathcal{L}_1	Autonomy Mode	No Fig.
Pitch Destabilized (-16.4% static margin)	NDI+ \mathcal{L}_1	Autonomy Mode	Fig. 8, 9, 10
Roll Destabilized	NDI+ \mathcal{L}_1	Autonomy Mode	Fig. 11

A. Control Performance in Nominal Flight Conditions

In the first flight test experiment we examined the performance of the NDI controller with and without the \mathcal{L}_1 AC law. This was a qualitative comparison of the overall L2F system with and without adaptive element of the controller in nominal conditions when there are no or small modeling errors. Due to the different real-time modeling results that were available to the two controllers during their respective flights, a true apples-to-apples comparison of the control methods cannot be made with the collected data. Instead, a carefully designed future experiment would be required to judge the performance of each control methodology.

Without \mathcal{L}_1 compensation: The first flight of the experiment began with an R/C takeoff. The R/C pilot engaged the learning mode with the PTIs once reaching the appropriate altitude. Data were collected over three laps of a racetrack pattern, two near the target speed and one at a lower speed. The PTIs were disengaged and the pilot performed pitch, yaw, and roll doublets. From here, the autopilot with only the NDI controller (i.e., no adaptive component) was given the control authority. There were no significant transients during this handoff. The vehicle began following the waypoints in the racetrack pattern. After demonstrating sufficient control of the vehicle, the autopilot was disengaged and reengaged multiple times at various headings, bank angles, and speeds. The results were similar for each case. Finally, the vehicle was returned to R/C mode for landing.

For one of the passes, Figure 5 shows the system states in blue and the commanded values in red. Good tracking performance can be seen in both the roll and pitch channels, but the yaw channel shows a small bias, especially in the sideslip angle β . Note that here, and throughout, we report the vehicle speed as the calibrated airspeed V_{CAS} .

With \mathcal{L}_1 compensation: The second flight of the experiment was similar to the first. It began with an R/C takeoff, after which, the learning mode was engaged with PTIs for a lap at three different speeds, low, medium and high. After disabling the PTIs, pitch, yaw, and roll doublets were performed. This time, when autopilot was engaged, the controller became the NDI controller with the adaptive component. Once again there were no significant transients after the switch, and the vehicle began navigating as expected. The tracking performance overall was good, as seen in Figure 6, but a small offset can be seen in the pitch channel, e.g. θ and q . Similar results were obtained after engaging the autopilot in

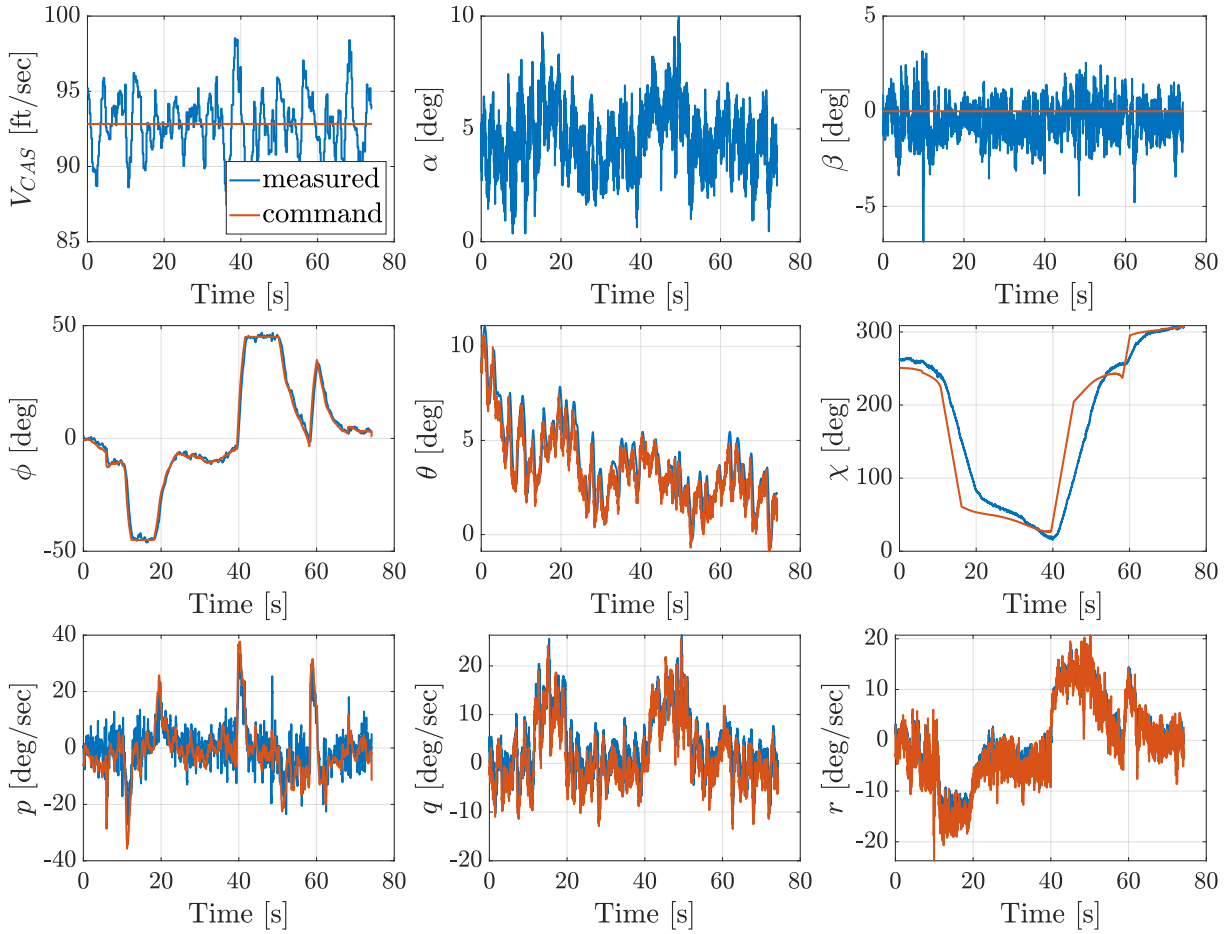


Fig. 5 Non-adaptive NDI Controller - Nominal Trajectory

numerous other conditions, including various roll and heading angles, at a high altitude with slow speed, and a low altitude with high speed. Satisfied with the performance, the pilot regained control with the PTIs active and flew several approaches to stall. The autopilot was then engaged after an approach to stall. As shown in Figure 7, it recovered the vehicle as expected and navigated to its waypoints. Finally, the vehicle was returned to R/C mode for landing.

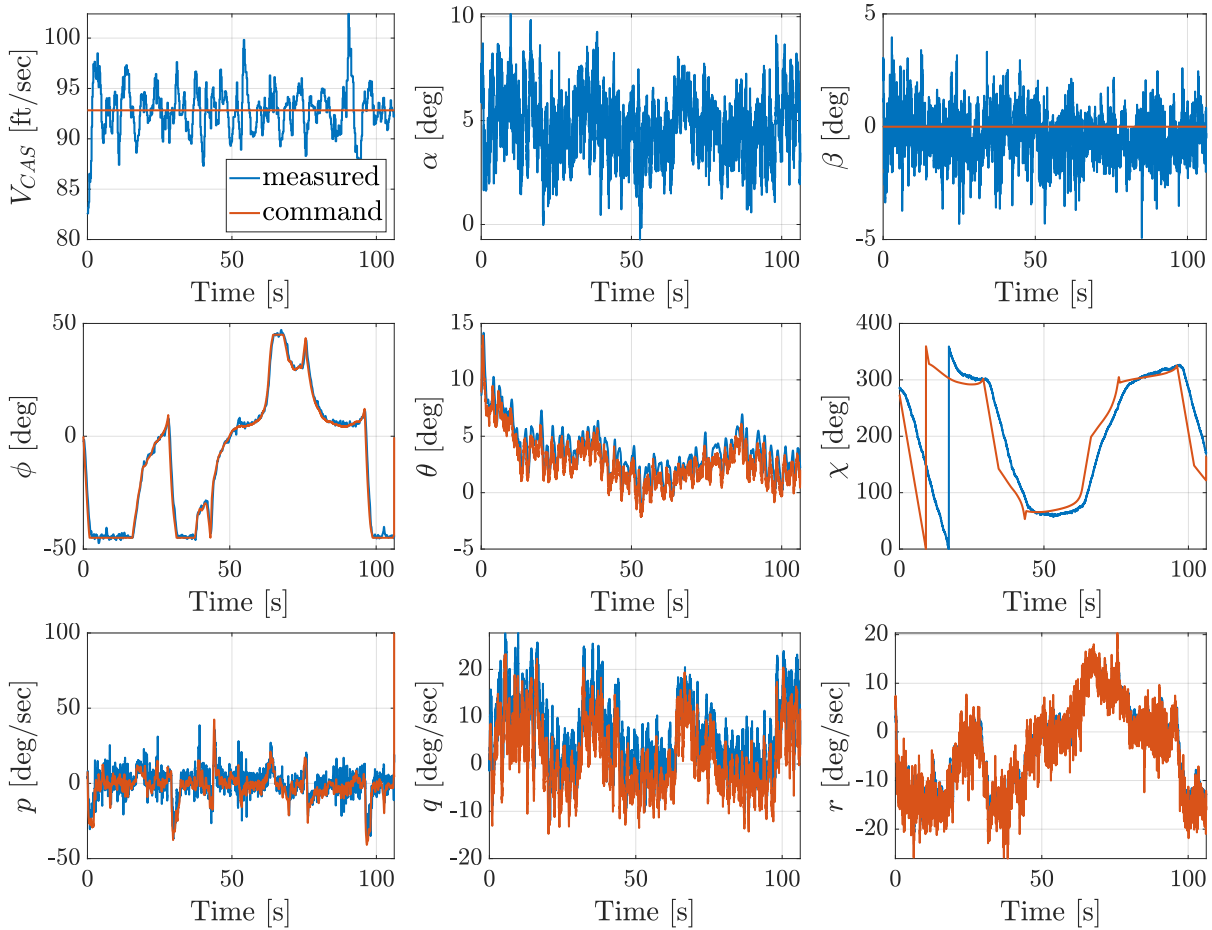


Fig. 6 Adaptive NDI Controller - Nominal Trajectory

B. Control and Learning Performance in the Presence of a Poor Initial Model

In the previous flights, the vehicle was manually flown through the learning phase while the PTIs were active. This was possible because the vehicle was stable and capable of being flown without the assistance of automatic control. In the next series of flights, a hidden feedback loop was used to destabilize the vehicle, in the pitch axis for the first two flights and in the roll axis for the third flight. This feedback loop was enabled at the same time as the learning mode and made the vehicle difficult to pilot manually. Thus in these flights, the learning and control occurred simultaneously. In addition, the actual (unstable) dynamics were significantly different from the (stable) initial model used by the controller, which is perfect for testing the capability of the L2F system. [In practice, the modeling inaccuracy may](#)

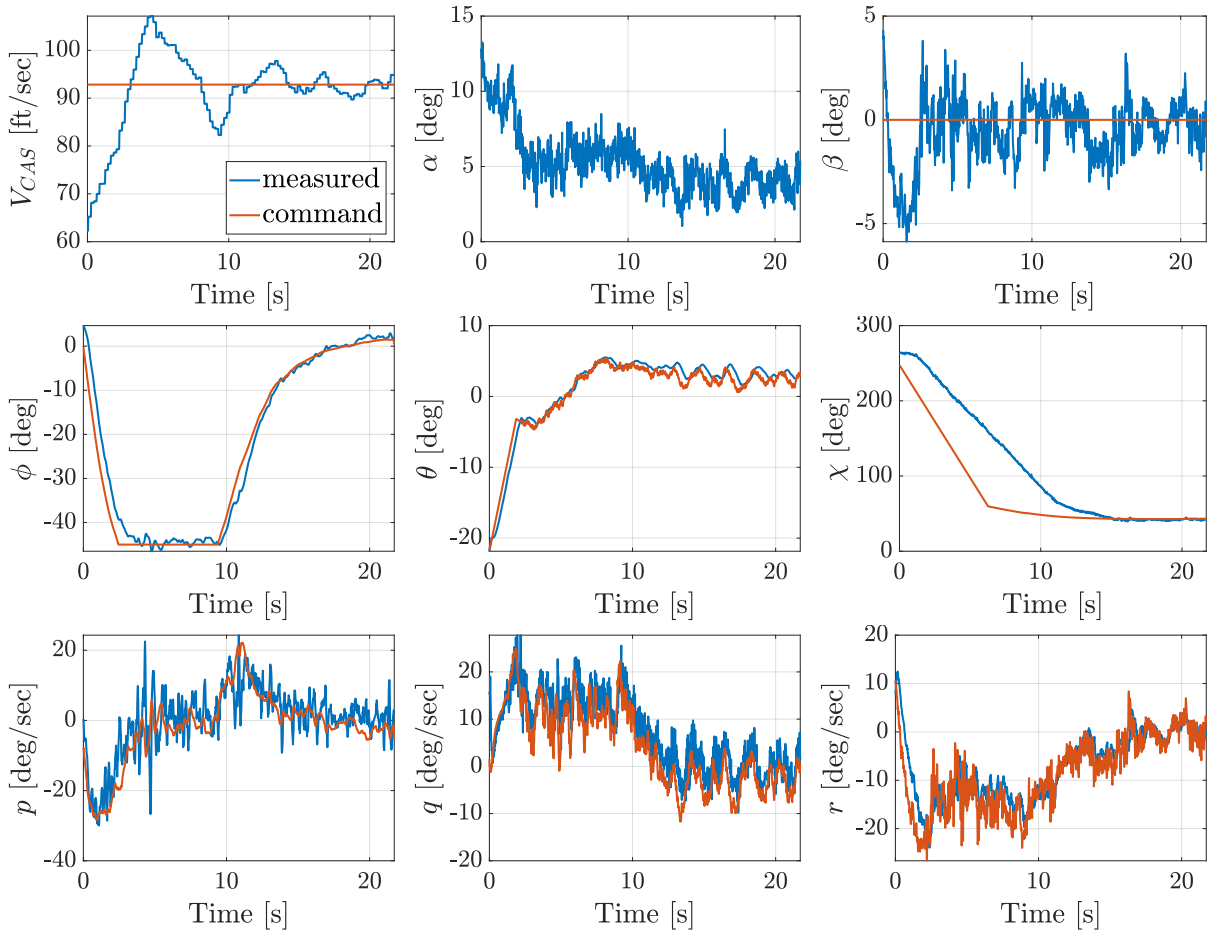


Fig. 7 Adaptive NDI Controller - Approach to Stall

contain unmatched and/or nonlinear uncertainties that were not considered in the theory of Section III; as a result, we may not obtain the same performance in the flight test results as in the theory.

The first flight in this series of experiments used angle-of-attack feedback to the left elevator to produce roughly a -10% static margin. The controller and modeling were not directly aware of this feedback and were blind to what the left elevator was doing; the modeling had no measurement of the left elevator's actions, and the controller could not send commands to the left elevator. The flight began with a piloted take off in the nominal configuration (no destabilizing feedback). Upon reaching the appropriate test altitude, the autonomous mode was engaged. This activated the destabilizing feedback, the PTIs, and the autopilot. The vehicle was able to stabilize and navigate to the waypoints. The autonomous mode was disengaged and reengaged several times, resulting in similar behavior. Using large stick inputs to attempt to maintain straight and level flight, the R/C pilot was also able to fly the destabilized vehicle in pilot command pass-through mode, though moderate angle-of-attack and pitch-attitude excursions were observed. The vehicle was returned to its nominal configuration for approach and landing.

The second flight was similar to the first but with a larger pitch-axis instability, about -16.4% static margin[‡]. Once again, the nominal vehicle was taken up to a certain altitude before engaging the destabilizing feedback, PTIs, and autopilot. A couple of pitch oscillations occurred, but the vehicle recovered and navigated appropriately. Repeats produced similar results. Unlike in the first flight, when the R/C pilot flew the destabilized aircraft, he was barely able to control it, even with full or almost-full stick deflections. Approach and landing were once again performed in the nominal configuration.

Figure 8 shows plots of the autonomous system's tracking performance from the second flight. Some transitory errors occurred at the beginning of the flight. However, after a few seconds, the vehicle settled and began tracking its commands with improved accuracy. In Figure 9, the adaptive command is plotted along with the pitch angle. After the learning transient, the vehicle was able to update its knowledge of the vehicle and the magnitude of the adaptive command reduced. Despite the improved system knowledge, uncertainties and disturbances still exist, such as atmospheric disturbances, the PTIs, etc. Because of this, the adaptive command did not necessarily drop to zero after learning.

The pilot's flight data (red) and that of the autonomous system (blue) are overlaid in Fig. 10. When the vehicle was controlled by the pilot, despite large inputs (δ_e) executed by the pilot, large pitch oscillations (θ) still occurred. In comparison, in the autonomous control mode, the real-time modeling was able to quickly determine that the vehicle was unstable (positive C_{M_α} estimate), and the control law modified the control gain to compensate. It is important to note that there will always be an inherent delay within the modeling results because sufficient informative data must be collected before a good model can be obtained. While real-time modeling was gathering the data, the baseline controller

[‡]A video of one of the flight tests comparing the performance of the R/C pilot and the L2F autopilot for the destabilized vehicle is available at <https://www.youtube.com/watch?v=y601mwzHdOE>

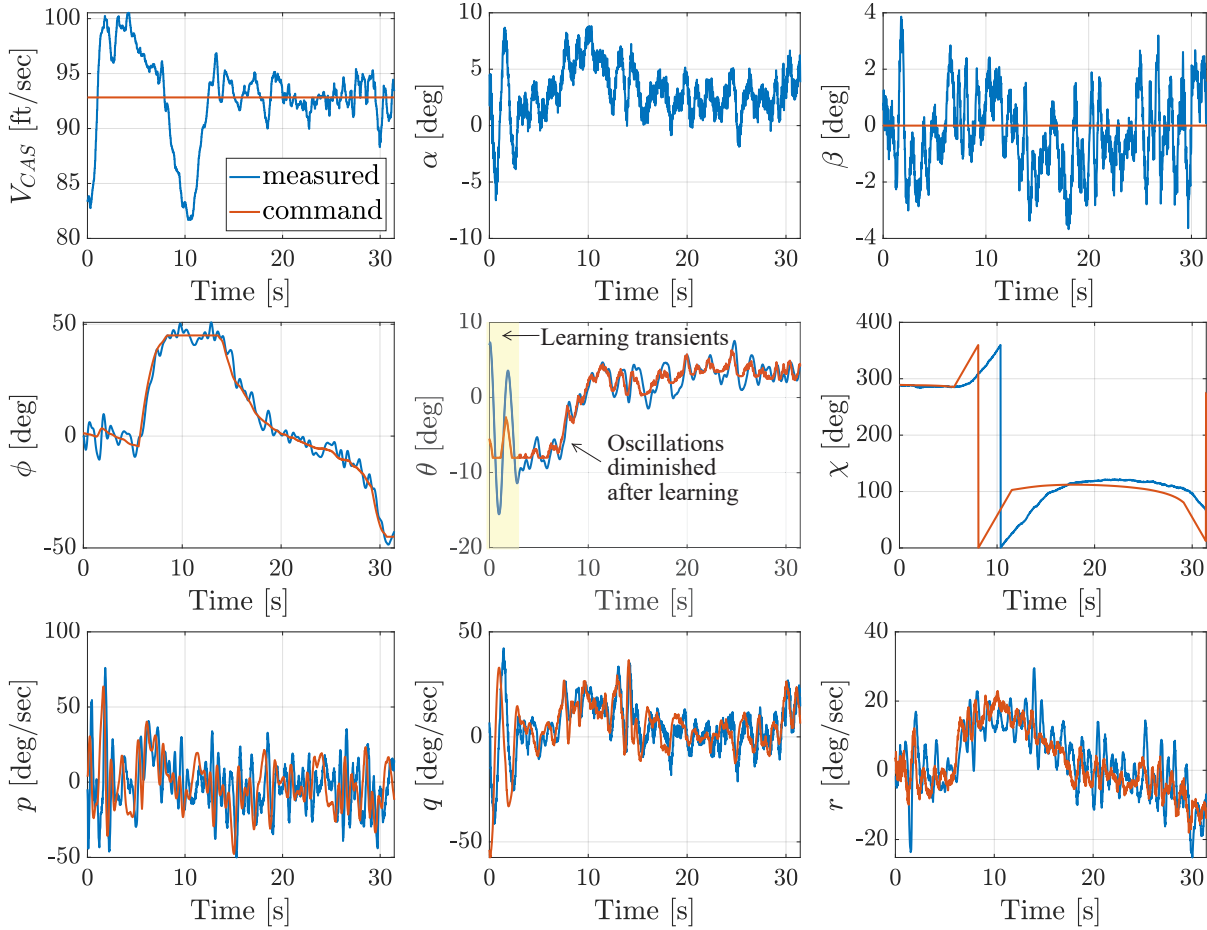


Fig. 8 Adaptive NDI Controller - Vehicle Pitch Axis Destabilization

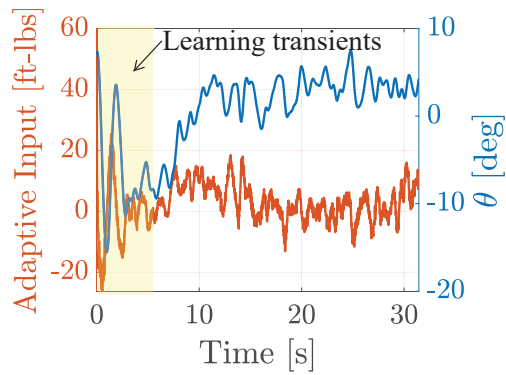


Fig. 9 Adaptive Input vs Pitch Angle for Pitch Axis Destabilization Case

was still operating based on a stable aircraft model (the initial guess) for the first few seconds. This is where adaptive compensation is critical for maintaining the stability of the vehicle. Note that due to use of the initial poor model, the actual uncertainties (i.e., the discrepancy between this model and the actual dynamics) most probably included unmatched and/or nonlinear uncertainties that are not considered in Section III. As a result, the tracking performance was not as good as claimed in Theorem 1. For instance, we can see a large pitch excursion in Fig. 10 at the beginning where the controller was struggling to achieve good tracking performance. Nevertheless, thanks to the adaptive compensation, the vehicle was able to maintain its stability, which is a major safety concern for an aerial vehicle. Once the system collected sufficient data to learn and update the model, the performance improved, resulting in relatively smooth, level flight.




Fig. 10 Overlay of Pilot Data with Autonomous Data

The third flight in this series used roll rate feedback to the inboard flaps in order to destabilize the roll axis. Initially, the target time to double was 0.5 seconds, but the feedback gain used was only able to render the vehicle neutrally stable. While the target level of instability was not achieved, it was sufficient to degrade the innate stability of the vehicle. Similar to the last two flights, the pilot first took off in the nominal configuration and then enabled the destabilizing feedback and PTIs and activated the autopilot mode once the vehicle reached the test altitude. Two engagements of the destabilizing feedback produced similar results with a few initial roll oscillations before navigating as expected. The R/C pilot also made two attempts to fly the destabilized vehicle. In the first, the pilot was able to stabilize the vehicle, but with some large roll oscillations. On the second, the pilot maintained better control of the vehicle but noted that it was difficult to fly. After reverting to the stable configuration, the pilot landed the vehicle.

Tracking data for this third flight can be seen in Fig. 11. Like the first two flights, there were large tracking errors at the beginning (though this time in the roll axis), but the L2F system quickly learned the new dynamics and the oscillations died down momentarily as the control laws were updated with more accurate models.

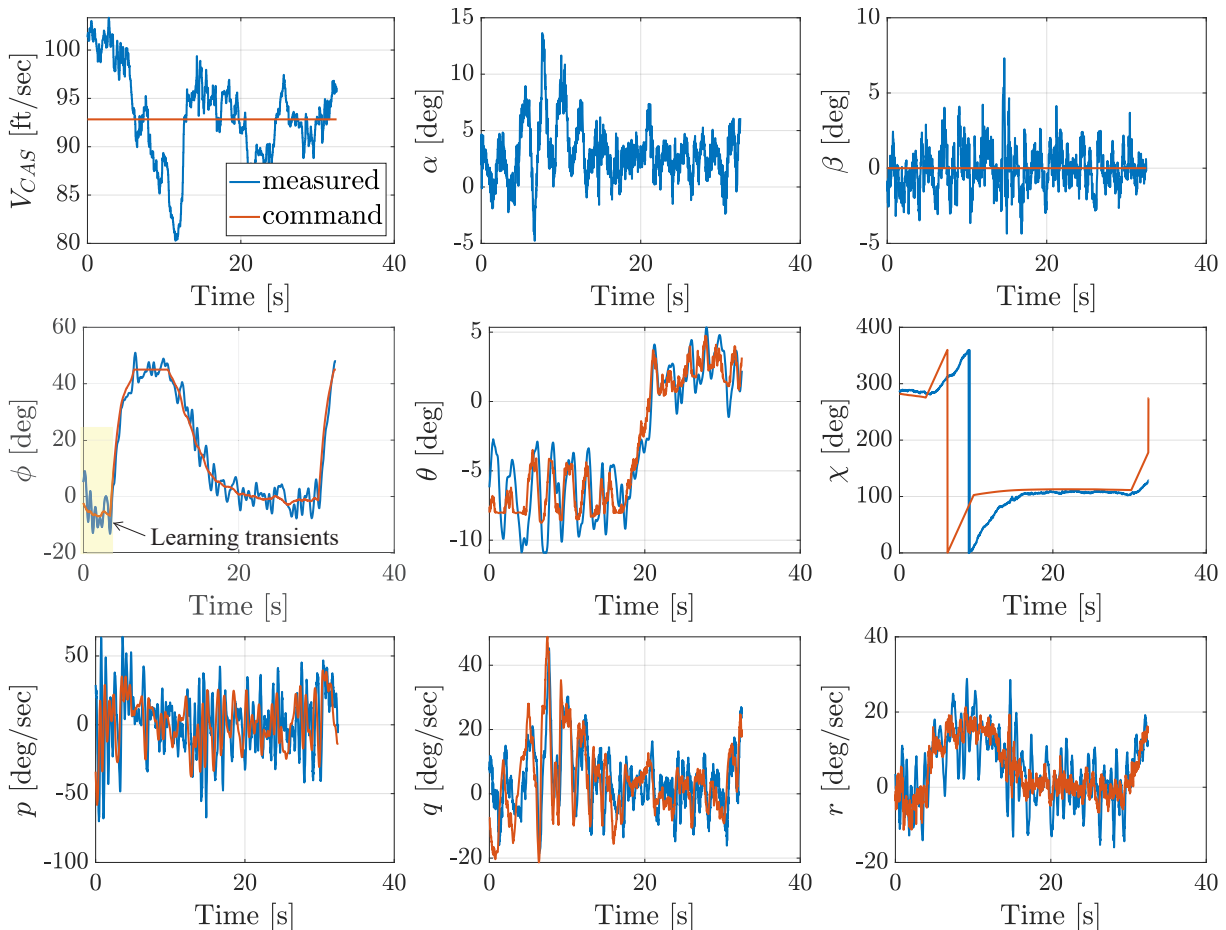


Fig. 11 Adaptive NDI Controller - Vehicle Roll Axis Destabilization

VI. Conclusions

This paper presents an \mathcal{L}_1 adaptive control (\mathcal{L}_1 AC) based scheme to ensure safe learning of aerial vehicle dynamics within the Learn-to-Fly (L2F) framework, which is envisioned to replace the traditional iterative development paradigm for aerial vehicles. At the core of the proposed scheme is an \mathcal{L}_1 AC architecture that provides stability and transient performance guarantees for a switched linear reference system subject to unknown time-varying parameters and disturbances. The paper also includes the controller design for the L2F system and how the \mathcal{L}_1 AC law is designed and integrated into the L2F system. The proposed scheme was validated by flight tests on an unmanned aerial vehicle.

Future work includes extension of the \mathcal{L}_1 AC architecture for switched systems subject to more general uncertainties (e.g., state-dependent and unmatched [20] uncertainties).

Acknowledgments

The authors acknowledge the generous long-term support from AFOSR and NASA for research on \mathcal{L}_1 adaptive control.

References

- [1] Heim, E. H., Viken, E., Brandon, J. M., and Croom, M. A., “NASA’s Learn-to-Fly Project Overview,” *Proceedings of AIAA Atmospheric Flight Mechanics Conference*, Atlanta, Georgia, 2018. doi:<https://doi.org/10.2514/6.2018-3307>.
- [2] Krstić, M., Kanellakopoulos, I., and Kokotović, P. V., *Nonlinear and Adaptive Control Design*, John Wiley & Sons, New York, NY, 1995.
- [3] Ioannou, P. A., and Sun, J., *Robust Adaptive Control*, Dover Publications, Inc., Mineola, NY, 2012.
- [4] Hovakimyan, N., and Cao, C., *\mathcal{L}_1 Adaptive Control Theory: Guaranteed Robustness with Fast Adaptation*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 2010. doi:<https://doi.org/10.1137/1.9780898719376>.
- [5] Cao, C., and Hovakimyan, N., “Design and Analysis of a Novel \mathcal{L}_1 Adaptive Control Architecture with Guaranteed Transient Performance,” *IEEE Transactions on Automatic Control*, Vol. 53, No. 2, 2008, pp. 586–591. doi:<https://doi.org/10.1109/ACC.2006.1657244>.
- [6] Hovakimyan, N., Cao, C., Kharisov, E., Xargay, E., and Gregory, I., “ \mathcal{L}_1 Adaptive Controller for Safety Critical Systems: Guaranteed Robustness with Fast Adaptation,” *IEEE Control Systems Magazine*, Vol. 55, 2011, pp. 54–104. doi:<https://doi.org/10.1109/mcs.2011.941961>.
- [7] Gregory, I., Cao, C., Xargay, E., Hovakimyan, N., and Zou, X., “ \mathcal{L}_1 Adaptive Control Design for NASA AirSTAR Flight Test Vehicle,” *Proceedings of AIAA Guidance, Navigation, and Control Conference*, Chicago, IL, 2009. doi:<https://doi.org/10.2514/6.2009-5738>, AIAA 2009-5738.
- [8] Gregory, I., Xargay, E., Cao, C., and Hovakimyan, N., “Flight Test of an \mathcal{L}_1 Adaptive Controller on the NASA AirSTAR Flight Test Vehicle,” *Proceedings of AIAA Guidance, Navigation, and Control Conference*, Toronto, Ontario, 2010. doi:<https://doi.org/10.2514/6.2010-8015>, AIAA 2010-8015.
- [9] Ackerman, K. A., Xargay, E., Choe, R., Hovakimyan, N., Cotting, M. C., Jeffrey, R. B., Blackstun, M. P., Fulkerson, T. P., Lau, T. R., and Stephens, S. S., “Evaluation of an \mathcal{L}_1 Adaptive Flight Control Law on Calspan’s Variable-Stability Learjet,” *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 4, 2017, pp. 1051–1060. doi:<https://doi.org/10.2514/1.G001730>.
- [10] Puig-Navarro, J., Ackerman, K., Hovakimyan, N., Cotting, M. C., Duke, D. J., Carrera, M. J., McCaskey, N. C., Esposito, D., Peterson, J. M., and Tellefsen, J. R., “An \mathcal{L}_1 Adaptive Stability Augmentation System Designed for MIL-HDBK-1797 Level 1 Flying Qualities,” *AIAA SciTech 2019 Forum*, San Diego, CA, 2019. doi:<https://doi.org/10.2514/6.2019-1083>, AIAA 2019-1083.

- [11] Chowdhary, G., and Johnson, E., “Concurrent Learning for Convergence in Adaptive Control Without Persistency of Excitation,” *49th IEEE Conference on Decision and Control*, IEEE, 2010, pp. 3674–3679. doi:<https://doi.org/10.1109/CDC.2010.5717148>.
- [12] Chowdhary, G., Yucelen, T., Mühlegg, M., and Johnson, E. N., “Concurrent Learning Adaptive Control of Linear Systems With Exponentially Convergent Bounds,” *International Journal of Adaptive Control and Signal Processing*, Vol. 27, No. 4, 2013, pp. 280–301. doi:<https://doi.org/10.1002/acs.2297>.
- [13] Sang, Q., and Tao, G., “Adaptive Control of Piecewise Linear Systems: The State Tracking Case,” *IEEE Transactions on Automatic Control*, Vol. 57, No. 2, 2011, pp. 522–528. doi:<https://doi.org/10.1109/TAC.2011.2164738>.
- [14] Xie, J., and Zhao, J., “Model Reference Adaptive Control for Switched LPV Systems and Its Application,” *IET Control Theory & Applications*, Vol. 10, No. 17, 2016, pp. 2204–2212. doi:<https://doi.org/10.1049/iet-cta.2015.1332>.
- [15] Wu, C., Zhao, J., and Sun, X.-M., “Adaptive Tracking Control for Uncertain Switched Systems Under Asynchronous Switching,” *International Journal of Robust and Nonlinear Control*, Vol. 25, No. 17, 2015, pp. 3457–3477. doi:<https://doi.org/10.1002/rnc.3275>.
- [16] Snyder, S., Zhao, P., and Hovakimyan, N., “ \mathcal{L}_1 Adaptive Control for Switching Reference Systems: Application to Flight Control,” *IFAC-PapersOnLine*, Vol. 52, No. 16, 2019, pp. 718–723. doi:<https://doi.org/10.1016/j.ifacol.2019.12.047>.
- [17] Morelli, E. A., “Practical Aspects of Real-Time Modeling for the Learn-to-Fly Concept,” *AIAA Atmospheric Flight Mechanics Conference*, Atlanta, GA, 2018. doi:<https://doi.org/10.2514/6.2018-3309>, dOI: 10.2514/6.2018-3309.
- [18] Foster, J. V., “Autonomous Guidance Algorithms for NASA Learn-to-Fly Technology Development,” *AIAA Atmospheric Flight Mechanics Conference*, Atlanta, GA, 2018. doi:<https://doi.org/10.2514/6.2018-3310>.
- [19] Liberzon, D., *Switching in Systems and Control*, Birkhäuser, Boston, MA, 2003.
- [20] Zhao, P., Snyder, S., Hovakimyan, N., and Cao, C., “Robust Adaptive Control of Linear Parameter-Varying Systems with Unmatched Uncertainties,” *arXiv:2010.04600*, 2020. URL <https://arxiv.org/abs/2010.04600>.
- [21] Zhao, P., Mao, Y., Tao, C., Hovakimyan, N., and Wang, X., “Adaptive Robust Quadratic Programs using Control Lyapunov and Barrier Functions,” *59th IEEE Conference on Decision and Control*, 2020, pp. 3353–3358. doi:<https://doi.org/10.1109/CDC42340.2020.9303829>.
- [22] Snell, S. A., “Nonlinear Dynamic-Inversion Flight Control of Supermaneuverable Aircraft,” Ph.D. thesis, University of Minnesota, October 1991.
- [23] Brandon, J. M., and Morelli, E. A., “Real-Time Global Nonlinear Aerodynamic Modeling from Flight Data,” *Journal of Aircraft*, Vol. 53, No. 5, 2016, pp. 729–740. doi:<https://doi.org/10.2514/1.C033133>.
- [24] Steven, B. L., and Lewis, F. L., *Aircraft Control and Simulation*, 2nd ed., John Wiley & Sons, Inc., 2003.
- [25] Riddick, S. E., Busan, R. C., Cox, D. E., and Laughter, S. A., “Learn-to-Fly Test Setup and Concept of Operations,” *AIAA Atmospheric Flight Mechanics Conference*, Atlanta, GA, 2018. doi:<https://doi.org/10.2514/6.2018-3308>.