Learn to Fly Test Setup and Concept of Operations

Stephen E. Riddick[1], Ronald C. Busan[2], David E. Cox[3], and Sean A. Laughter[4]
NASA Langley Research Center, Hampton, VA, 23681

The NASA Learn-to-Fly (L2F) project recently completed a series of flight demonstrations of its learning algorithm for flight control at Fort A. P. Hill in Virginia. This paper discusses the test setup and concept of operations (ConOps) used by the L2F team. Unmanned airframe demonstrators for testing the research algorithms included a modified commercial off-the-shelf subscale powered airplane, plus four gliders – two of which had an unconventional configuration and were fabricated using a "rapid" prototyping technique. Avionics system similarities and differences between the test aircraft are described, as well as ground testing in preparation for flight. The ConOps discussion includes the development of a tethered helium balloon drop launch technique for the glider demonstrators. This launch method was chosen for its potential to be inexpensive and allow for rapid turn-around for multiple glider launches – but it also presented challenges, such as balloon tether avoidance, high angle of attack, low dynamic pressure initial conditions, and susceptibility to winds. A remotely piloted approach employing high-end hobbyist radio controlled (R/C) hardware was used for the powered demonstrator. This approach accommodated the interaction between the R/C flight system and the research flight control computer, engaging the L2F algorithm at varying initial conditions and artificially reducing the aircraft stability to stress the algorithm.

I. Nomenclature

ADC	=	analog-to-digital converter
A/D	=	analog-to-digital converter
AGL	=	above ground level altitude (laser altimeter)
AoA	=	angle of attack
AoS	=	angle of sideslip
ARF	=	almost-ready-to-fly
ARMD	=	Aeronautics Research Mission Directorate
CAS	=	Convergent Aeronautics Solutions Project
CERTAIN	=	City Environment for Range Testing of Autonomous Integrated Navigation
CFD	=	computational fluid dynamics
Cm_α	=	pitch stability, coefficient of pitching moment derivative with respect to angle of attack
ConOps	=	concept of operations
COTS	=	commercial off the shelf
CPLD	=	complex programmable logic device
FAA	=	Federal Aviation Administration
FCC	=	flight control computer
FTS	=	flight termination system
GPS	=	Global Positioning System
IMU	=	inertial measurement unit
INS	=	inertial navigation system
k_α	=	effective proportional gain for the longitudinal axis
L2F	=	Learn-to-Fly
LiFe	=	lithium iron phosphate battery

[1] Research Engineer, Flight Dynamics Branch, MS 308, Member
[2] Senior Research Engineer, Flight Dynamics Branch, MS 308, Associate Fellow
[3] Senior Research Engineer, Dynamic Systems & Control Branch, MS 308
[4] Electrical Engineer, Aeronautics Systems Engineering Branch, MS 328

LiPoly	=	lithium polymer battery
NASA	=	National Aeronautics and Space Administration
OAT	=	outside air temperature
PCIe	=	Peripheral Component Interconnect Express
PTIs	=	Programmed Test Inputs
PPT2	=	precision pressure transducer
PWM	=	pulse width modulation (servo command signal)
R/C	=	radio controlled
RPM	=	rotations per minute (motor speed)
SBUS	=	Futaba proprietary servo command bus system
TACP	=	Transformative Aeronautics Concepts Program
TRL	=	Technology Readiness Level
USB	=	universal serial bus
V	=	volt
v	=	volt

II. Introduction

THE Learn-To-Fly (L2F) concept is a subproject under the Convergent Aeronautics Solutions (CAS) Project under NASA's Transformative Aeronautics Concepts Program (TACP) [1-4]. The CAS project is an avenue to fund feasibility studies of low Technology Readiness Level (TRL) concepts with high-potential payoffs for large-scale aeronautics challenges. Specifically, L2F is a paradigm change on how the process of aircraft control development is conducted. The conventional iterative loop consists of testing, simulation development, control design, and flight test. Instead of the traditional iterative process, the modeling, validation, and controls development are all done concurrently and autonomously with near real-time modeling of vehicle aerodynamics and other characteristics coupled with a novel learning control system. The real-time modeling and learning control systems that comprise L2F have the potential to radically change and shorten the flight development testing process, and improve safety due to real-time learning of the airplane itself. The complexity of many interacting effects (multiple distributed electric engines, flexible aircraft, innovative controllers, etc.) will tax even the ability to fully characterize aerodynamic and interaction effects with conventional test techniques. L2F could enable the complex concepts to rapidly get to flight and perform to the design's maximum capabilities in a potentially faster and more reliable approach than current methods.

The L2F approach combines real-time modeling of aircraft with adaptive controls to rapidly converge on an optimized control solution, starting with only a very rudimentary assumption of flight parameters. To achieve the real-time aerodynamic modeling of the system, the algorithm uses recursive multivariate orthogonal functions and programmed test inputs (PTIs) on each control surface and observes aircraft response. The PTIs are preprogramed multi-sine inputs that are superimposed on surface commands. Because they are orthogonal, the PTIs can be applied to all control surfaces simultaneously. The algorithm then forms a model from the observation and then uses the model to adjust the control part of the algorithm. As the flight progresses, the PTIs continue to refine the model, from which the controls are improved. The modeling function of the algorithm is described in Ref. [2] and [5], while the learning control architecture is illustrated in Ref. [3]. Reference [4] describes the guidance algorithm, which serves as an "executive" function.

To achieve the L2F objectives, the project chose three types of vehicles. The first consisted of foam gliders used as proof-of-concept for a balloon drop system used to launch glider models. The second design was an unusual aircraft configuration with the purpose of presenting a significant challenge for the algorithm. The third was a conventional aerobatic configuration that could vary its stability characteristics to simulate unstable aircraft. No conventional experimental aerodynamic knowledge, such as captive static and dynamic wind tunnel test results, was used to guide the process before flight. The initial assumption for the L2F algorithm was based on rudimentary vortex lattice codes [1]. This was done to remove any temptation to tune the algorithm before flight, thereby forcing it to learn on its own. This also meant that even though numerous simulations were performed for the algorithm development, the simulation intentionally lacked an accurate aerodynamic model.

This paper will present the concept of operations and describe the balloon launch system (Section III) used to launch the L2F gliders, as well as the challenges inherent in that approach. The gliders are described in Sections IV and V. The powered aircraft description and ConOps is found in Section VI. The avionics used to accomplish the L2F mission is presented in Section VII. Although a more comprehensive discussion on project results can be found in Refs. [2-4], Section VIII provides a sample of those results.

III. Balloon Launch System

The choice of a tethered balloon as a launch platform was the result of practical and programmatic considerations. In an attempt to simplify the L2F algorithm, a glider concept was chosen for the initial flight demonstrators so that propulsion effects could be ignored. Several options were considered for getting the glider to an altitude high enough to test the L2F algorithms and systems. Using a helicopter as a drop ship would be relatively expensive for multiple drops – as well as impose more extensive safety and operational overhead associated with having manned aircraft as part of the test program. Release from a powered remotely piloted vehicle was briefly considered. However, it would be impossible to reasonably predict the possibility of re-contact due to the intended lack of high fidelity data for use in analyzing flight dynamics immediately after launch. Therefore, a tethered helium aerostat was chosen as a relatively inexpensive carrier that had the potential of permitting multiple launches per day.

A. Balloon Description

Two different size balloons were used during the L2F project. The first held 1000 ft^3 of helium and was rated for 45 lbs of lift, while the second held 2000 ft^3 with a rated lift of 90 lbs. Both balloons were similar in design and construction. Each consisted of an oblate spheroid single-ply helium-filled urethane bag with a fabric sail attached to the lower half on one side. The design was called a "kytoon" or "kite balloon," because it obtained some of its lift dynamically as a heavier-than-air kite and the rest aerostatically as a lighter-than-air balloon. Kytoons have been used for civilian and military purposes for almost 100 years because they tended to be more stable in varying winds than ordinary tethered balloons or kites [6]. This configuration is illustrated in Fig. 1.

The fabric sail caused the balloon to weathervane, assuring a consistent alignment of the balloon relative to the wind direction. Although the fabric sail was not as large as on some kytoons, the aerodynamic lift was still calculated to be greater than the aerostatic lift from the helium for moderate or higher wind speeds. Because the L2F carrier and airframe cross-sectional area exposed to the wind was significantly greater than the packages typically carried by these balloons, drag was proportionally much higher. Therefore, the stability in winds was significantly less than what the manufacturer quoted in their literature – although still better than for a balloon with no fabric sail (i.e., one relying solely on aerostatic lift).

In the case of a tether failure, the balloon was equipped with an FAA-approved flight termination system (FTS). This system used a commercial off the shelf (COTS) radio controlled (R/C) system that controls a heated wire circuit that melts a hole in the balloon urethane skin when commanded. This would release helium causing the balloon to descend – with the rate of descent being dependent on how long the FTS was commanded and how large a hole was produced.

The smaller balloon was used for the initial L2F gliders described in Section IV, while the larger balloon was used for the Woodstock gliders described in Section V. The Woodstock gliders were the primary focus of the glider phase of the L2F project; therefore, this paper is focused on that particular balloon system.

B. Suspension Line Setup

Three load-bearing suspension lines (one from each attachment point) connected to the apex of the carrier structure and the tether (Fig. 1). Each balloon attachment point also had two additional supporting suspension lines that connected to the other two corners of the carrier structure. These lines oriented the structure to set the glider release angle and create a tensioned truss-like arrangement that kept the carrier and airframe acting as if they were rigidly attached to the balloon overhead.

C. Tether and Winch System

A single tether line (Fig. 1) anchored the helium balloon to the support trailer. For the larger balloon the tether was 1/8" diameter line made from Spectra fibers. The tether had a 3000 lb. minimum breaking strength and weighed 4.5 lbs. per 1000 feet of length.

An electric winch (Fig. 2) was used to spool the tether in or out. For the larger balloon the winch used a 4 horsepower electric motor run off of two 12V marine batteries. A variable speed control allowed the tether to be deployed at up to 160 feet per minute. During retraction and retrieval of the balloon, the motor had to overcome the load associated with the balloon lift and drag from any wind. When the destination altitude was several thousand feet, this could deplete the batteries – so the system had a built-in charger that can be connected to an on-sight generator to extend the time that the winch can be run continuously.

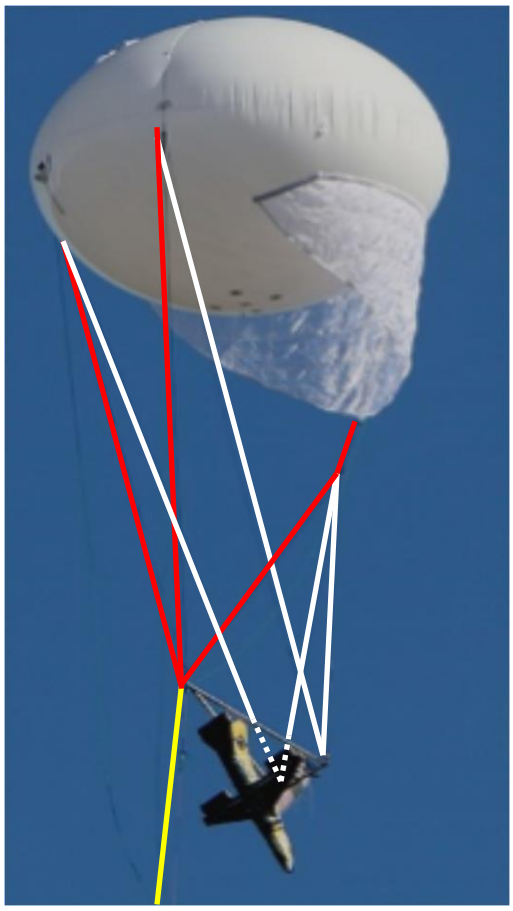

Fig. 1 Line configuration. Tether (yellow), load-bearing suspension lines (red), supporting suspension lines (white).

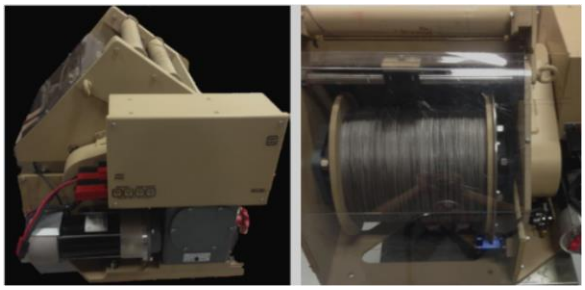

Fig. 2 Winch for 2000 cubic foot helium balloon.

D. Balloon and Glider Concept of Operations

The aerostat system had three attachment points located approximately 120 degrees apart around the equator of the balloon. A metal ring at each attachment point was used to secure one handling line and several suspension lines. The suspension lines then attached to the carrier structure, which held the flight article and the tether connection. Once the glider preflight checklist was completed, the balloon was positioned over the glider using the handling lines (not visible in Fig. 1) as the flight article was attached to the carrier glider release mechanism. Once the glider was connected and a flight authorization has been given, the balloon FTS unit was armed and the handling lines are secured above the carrier structure to prevent fouling the glider during ascent or after launch. At this point, the balloon was being held by the tether alone.

The tether was then spooled out with an electric winch to a predetermined flight initiation altitude. This altitude was determined by either referencing flags attached to the tether at known intervals, or through an independent altitude sensor and data link from the balloon.

Once the target altitude was reached, the remote pilot would command the release of the glider. Under pilot observation, the glider would perform a pullout maneuver before starting PTIs. It would then execute its learning process while performing waypoint navigation. At a predetermined altitude, it would go into a holding pattern, then implement a landing pattern, flaring for a touch down, all under full autonomy [4]. However, the glider was considered expendable and not expected to survive the landing due to estimated approach speed. The pilot has the ability to assume control when confronted with a safety issue, a range boundary violation, or a L2F control system failure. Landing could also be controlled directly by the remote pilot, but given the intentional lack of preflight analysis, the airframe was never guaranteed to be flyable. The entire operation was overseen by a range safety officer and coordinated by a test director. The balloon flight termination operator and the winch operator were in constant radio communication with the rest of the team. These ConOps are visualized in Fig. 3.

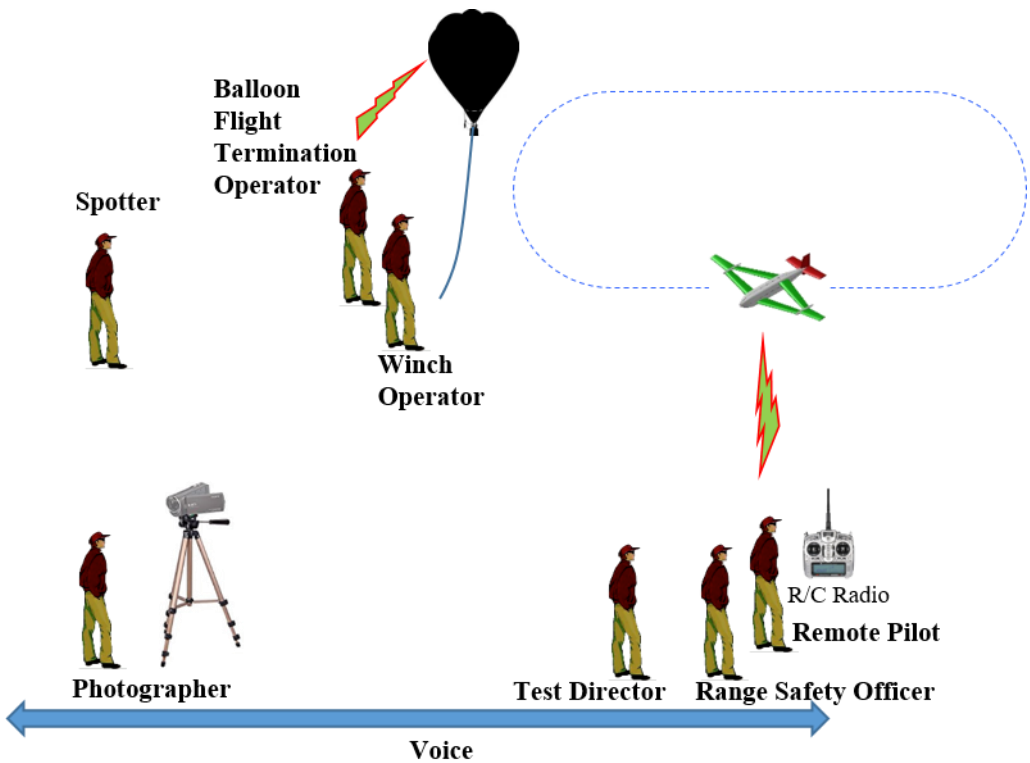

Fig. 3 Balloon ConOps Visualization.

E. Balloon Launch Considerations and Challenges

The first major consideration is preventing flight article contact with the tether after release. To mitigate this, the glider is oriented pointing away from the tether in hopes that it would fly away from the hazard. This is illustrated in Fig. 1. The drawback is that this creates an extremely high angle of attack for the glider immediately after release, especially in the presence of wind. The wind pushes the balloon away from the winch station, so the tether is upwind of the balloon. The glider is pointed away from the tether, forcing a tailwind initial condition.

Airframe release attitude was another major consideration. Early tests with foam gliders showed that if the airframe nose was pointed straight down to align it with the initial flight path, a tail wind would rotate the glider past a 90-degree pitch attitude immediately after release, worsening the severity of the pullout maneuver. If the flight article was released in a level orientation, an initial angle attack was guaranteed to be a minimum of 90 deg. Also, once released, the aircraft had to build airspeed and perform a pull-up maneuver before it could assume any quasi-straight-and-level test condition. The project ultimately chose a release angle of approximately 60 degrees nose down for the final drops.

A lesson learned was that determining release altitude by using indicators at premeasured locations on a tether line is less than ideal. Buoyancy forces are reduced as altitude increases due to the decrease in air density. Therefore, the tendency of the balloon to travel downrange increases with altitude, especially in the presence of wind. If the buoyancy margin above the payload weight is small, a point can be reached where the release of more tether line increases the horizontal distance from the winch but does not increase altitude. A barometric measurement device with a data downlink was shown to be a more accurate way to determine release altitude.

Overall, the tethered balloon system can be a very efficient method of drop testing if the aircraft is recovered without damage and able to be reused. The balloon system could also be efficient if multiple expendable airframes are readied. It was determined that having only the two Woodstock airframes (described in Section V) was not sufficient to make efficient use of the balloon drop system.

Fig. 4 Balloon system being prepared for deployment.

IV. Preliminary Glider Configurations

Two different foam airframes were used as proof-of-concept in the development of the balloon drop technique (called “Foamy” and “Super Guppy Foamy”). Both of them were modified MiG-27 target drones from RS Systems (model AN/FQM-117B). Due to the low drop altitudes of initial testing and the associated low speed flight, the wingtips of these airframes were extended. Also, the ailerons and elevators were enlarged to boost control power. Reference dimensions and weights of the two glider configurations are shown in Table 1. Neither airframe had landing gear, so landings consisted of a belly skid on grass. Both vehicles were dropped from planned altitudes of 400 ft. from a tethered balloon at the North Range of the City Environment for Range Testing of Autonomous Integrated Navigation (CERTAIN) test range at NASA Langley Research Center. Results of these tests can be found in Ref. [1].

Table 1 Reference dimensions and weight of initial L2F gliders.

	"Foamy"	"Super Guppy Foamy"
Wing area ,sq. ft.	7.56	7.56
Mean aerodynamic chord, ft.	1.8	1.8
Wing Span, ft.	4.76	4.76
Weight, lbs.	7.5	15.4

A. Foamy Description

The first airframe (Fig. 5), dubbed "Foamy" weighed approximately 7.5 lbs. and carried a rudimentary research avionics package. The associated sensor suite was minimal and only measured GPS position along with inertial parameters and static pressure. This flight article was used to test various release attitudes and provided a proof-of-concept for the balloon drop technique. The foam glider flights showed promise, but a faster research computer with a more sophisticated sensor suite was needed.

Fig. 5 "Foamy" airframe during a control surface test.

B. Super Guppy Foamy Description

The second airframe (Fig. 6), called "Super Guppy Foamy", was further modified to expand the internal fuselage volume. This was done by making a conformal pod out of a spare fuselage and was required to accommodate the production L2F avionics, including power distribution, switcher boxes, and a larger flight computer. Added sensors included an air data boom and control surface encoders. The wing was strengthened with two spars to handle the added weight. The final gross weight for this airframe was 15.4 lbs, more than twice the previous foam model. It was intended as a checkout for the avionics package, but its only flight resulted in a crash that destroyed the airframe.

Fig. 6 "Super Guppy Foamy" airframe incorporates a belly extension to house the production Learn-to-Fly avionics.

V. "Woodstock" Glider

The L2F project was tasked with testing the feasibility of using advanced modeling and adaptive flight controls to achieve flight without the aid of the conventional costly and time-consuming iterative loop of: testing; simulation development; control design; flight test; repeat. To accomplish this with any conviction, a novel design (i.e., one in which flight dynamics were not obvious to an experienced engineer) was required. Also, once the aircraft was designed, aerodynamic test results from high-fidelity ground-based methods (wind tunnels, CFD, etc.). would not be used. Instead, only rudimentary vortex lattice codes. These decisions were made to prevent the control designers from "cheating" by designing to known rules of thumb or tuning to an accurate simulation model. The final glider design was dubbed "Woodstock." Two airframes were constructed and supplied with identical avionics suites. Because aircraft flight characteristics were never tested and the ConOps involved an autonomous landing, these vehicles were considered expendable. With a release altitude between 2000 and 3000 feet, each airframe made one flight attempt at Finnegan Field within Fort A.P. Hill. Results from those flights are found in Refs. [1-4].

A. Woodstock Description

The Woodstock airframe (Fig. 7) consisted of a rectangular fuselage and a diamond wing with wingtip extensions. Eight control surfaces were placed on the trailing edges of the diamond while another two were placed on the tip extensions. The aft fuselage was designed with a dual-axis gimbal joint placed just forward of the empennage. This allowed the entire empennage to move as both an all-moving vertical and horizontal tail. The wingspan was 91.2 in. and the weight was 27.5 lbs. The primary avionics consisted of a high-end COTS R/C receiver and was powered by a single 4 amp-hour lithium iron phosphate (LiFe) battery. The research system was powered by a 5 amp-hour lithium polymer (LiPoly) battery. These components are listed in Table 2. The avionics and research systems are described in section VII.

Fig. 7 Woodstock airframes.

Table 2 Woodstock at a glance.

Airframe	Custom 3D printed diamond wing with gimballed tail configuration
Wing area	6.94 sq. ft.
Mean aerodynamic chord	1.67 ft.
Wing Span	7.6 ft.
Weight	28.35 lbs.
Airspeed	20-80 kts.
Endurance	~5 min.
Primary avionics system	Spektrum DX18 (2.4 GHz)
Avionics battery	1x4,000 mAh 2-cell LiFe (6.6v nominal)
Research battery	1x5,000 mAh 18.5v 5-cell LiPoly (18.5v nominal)

B. Woodstock Construction

Woodstock was constructed using an experimental approach for the rapid design and fabrication of aircraft. The entire airframe was 3D printed from nylon. This approach was expected to save both time and cost, but it had numerous unforeseen shortcomings, which caused significant delays in airframe delivery. The nylon proved to be too flexible, so the wings had to be re-designed to accommodate carbon fiber spars. The remaining flexibility caused the control surfaces to bind, so they had to be fitted by hand, increasing fabrication time. The tail had to be reinforced with aluminum to handle the loads from the all-moving empennage. These changes to the original design caused the final airframe weight to be about twice the original estimate. In hindsight, a more traditional construction approach would have been faster, lighter, and more durable. However, the design and fabrication team gained valuable input for improvements to their experimental approach.

C. Woodstock Fate

Unfortunately, the first Woodstock airframe re-contacted the release system, breaking the servo arm that controlled the pitch axis of the empennage. The resulting flight consisted of an inverted spiral and a crash. Despite the lack of control, the L2F algorithm did obtain a modeling estimate at an angle of attack of approximately -38 degrees. The second Woodstock airframe attempted a turn towards a waypoint while performing the pull-up maneuver. This initiated a spin, which the L2F algorithm arrested. During recovery, the airframe entered a spin in the opposite direction, that was never recovered. Feasibility of the L2F concept could not be assessed with either Woodstock flight.

VI. Powered Aircraft

The powered aircraft for Learn-to-Fly was an electric propeller-driven aircraft dubbed “E1.” It was controlled by a remote pilot who performed takeoffs and landings manually. Test points included assessing the L2F algorithm as a developmental autopilot, engaging the algorithm in off-nominal conditions, and stressing the algorithm by destabilizing the aircraft. Eleven flights were performed with this platform.

A. E1 Description

The airframe chosen for the Learn-to-Fly powered aircraft was an almost-ready-to-fly (ARF) kit of a 40% scale Extra 330 SC (Fig. 8). The original kit had an aileron configuration running the entire trailing edge of the wing. The kit was modified by cutting those surfaces in half and using the inboard half as a flap with positive and negative deflection capability. An electric propulsion system was chosen to avoid the complexity of dealing with fuel. The ability to power off the motor and restart it in flight considered to be another advantage. The propeller was driven by a 15 kW motor powered by 6 LiPoly 8Ah batteries. They were arranged in 2 parallel sets of 3 wired in series, giving a nominal system voltage of 55.5v. The primary avionics consisted of a high-end COTS R/C receiver and was powered by a pair of redundant 7.8 Ah LiPoly batteries. The research system was powered by a 5 Ah LiPoly battery. Table 3 provides a list of these components. The avionics and research systems are described in Section VII.

Fig. 8 E1 during R/C range tests.

Table 3 E1 at a glance.

Airframe	40%-scale Extra 330 SC
Wing area	19.26 sq. ft.
Mean aerodynamic chord	1.97 ft.
Wing Span	10.17 ft.
Weight	61.45 lb.
Airspeed	35-88 kts.
Endurance	~15 min.
Primary avionics system	Futaba 18MZ (2.4 GHz)
Avionics battery	2x7,800 mAh 2-cell LiPoly (7.4v nominal)
Motor	Hacker 200-8 (15 kW)
Motor battery	2x3x8000 mAh 18.5v v 5-cell LiPoly (55.5v nominal)
Research battery	1x5,000 mAh 18.5v 5-cell LiPoly (18.5v nominal)

B. E1 Concept of Operations

E1 was flown using a remote pilot on a COTS R/C system. The pilot would take off and land in a safety mode that ignored the research system. The research system had two modes. The first was a pilot command pass-through with the option to overlay inputs from the research system (PTIs, stability augmentation, etc.). The second was a fully autonomous mode that gave the research system full control of the aircraft. At any time, the pilot could revert to safety mode, bypassing the research system altogether.

The pilot would take off, fly to a predetermined initial condition, and engage the research system. After the test point was completed, the pilot would revert back to safety mode and maneuver to the initial condition for the next test point and repeat. When the test points were completed or a safety time limit was reached, the pilot would revert to the

safety mode and land the aircraft. A range safety officer oversaw the operation, while a test director coordinated the flight plan with the pilot. This is illustrated in Fig. 9.

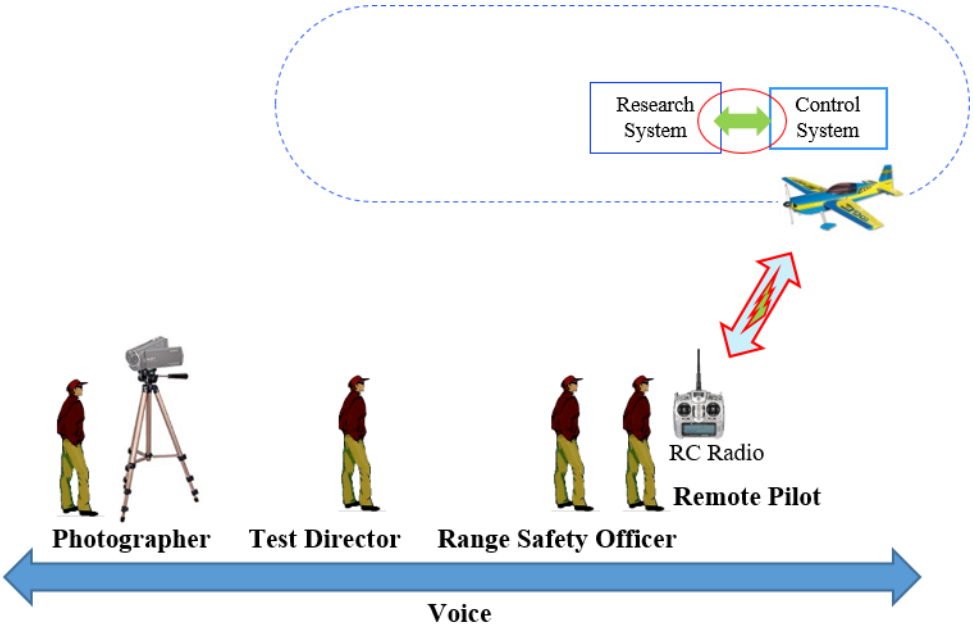

Fig. 9 E1 ConOps.

C. Engagement Scenarios

One goal for the E1 test flights was to evaluate the feasibility of using the L2F algorithm as a developmental autopilot. This was done in several ways. The first was a phased process. After E1 was put in the proper flight condition, the pilot enabled the modeling portion of the flight code. Programmed Test Inputs (PTIs) overlaid pilot commands, allowing the L2F software to define and refine an initial aerodynamic model. After a set time, the adaptive control and guidance algorithms were engaged, putting the aircraft in a fully autonomous mode. During subsequent engagements, the time between phases was shortened to determine a minimum time needed for modeling to operate before determining an adequate model.

The second evaluation technique was meant to simulate a disoriented pilot using the autopilot as a last resort. For this evaluation, the modeling algorithm was allowed to remember the model it refined on previous engagements. The full autopilot was initially engaged on course to confirm it maintained that course. Following a buildup process, it was then engaged with the airplane purposely off heading, forcing a course correction. Later, the autopilot was engaged off attitude, including high pitch and bank angles. Finally, it was engaged with heading, attitude, altitude and speed excursions to see if it recovered to the expected flight condition. The last evaluation technique involves varying the stability, which is discussed in the next section.

D. Variable Stability

The ability of the L2F algorithm to handle degraded stability was evaluated by synthetically altering the aircraft stability with the use of dedicated flight controls. The L2F software had no use or knowledge of the existence of these "instability surfaces." To destabilize the roll axis, the flaps were tied to roll rate. When a roll rate was detected, the flaps moved in such a way to increase that roll rate (Fig. 10). The left elevator was tied to angle of attack to destabilize pitch. As angle of attack increased, the surface would move in an attempt to increase angle of attack further. Both axes were destabilized via a proportional gain that could be adjusted between flights. This degraded stability mode could be used both under pilot control and fully autonomous control. As stated previously, takeoffs and landings were always performed in a safety mode that disabled all research software and gave the pilot complete access to all control surfaces. This technique allowed for a very rapid buildup process of increasing aircraft instability. For example, the test team was able to perform three of these variable stability flights inside of an hour.

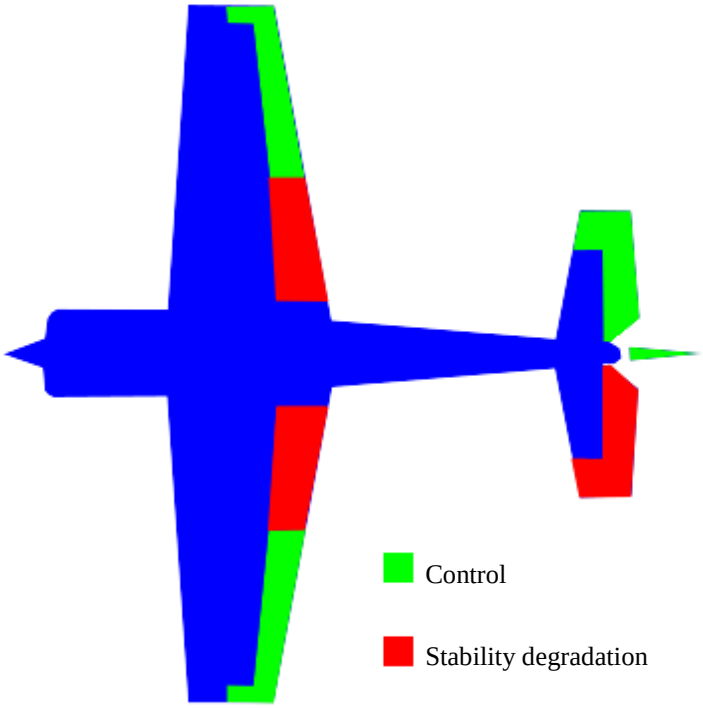

Fig. 10 E1 degraded stability control allocation.

VII. Avionics

The philosophy behind the avionics design was to have a high-end commercial off-the-shelf (COTS) R/C system serve as the primary flight avionics. It served as the sole link between the pilot and the aircraft, as well as being responsible for the lost-link failsafe profile. The research control system consisted of a mix of COTS and custom hardware, as well as custom software. The research system could only operate under certain conditions, as determined by the primary flight avionics, but required a deliberate selection by the external pilot via a permissive switch. Note that the avionics packages vary slightly between the two aircraft types because they were designed and built by two different teams.

A. Woodstock Flight Avionics

The Woodstock avionics (Fig. 11) consisted of COTS components (Spektrum R/C transmitter, receiver, batteries, and servos) along with custom in-house components including a power distribution unit and a pulse width modulated (PWM) servo command selector switch. The command switch between the receiver and the research system was a custom component designed and built in house. It was controlled by the R/C receiver signal and selected whether the avionics system or the research system was in control of flight surfaces.

B. E1 Flight Avionics

The E1 avionics components (Fig. 12) are from Futaba. They include the servos, transmitter, and redundant dual receiver setup. The receiver outputs are fed into a PowerBox Champion, which combines the two receiver outputs (surface commands) and detects failures and signal dropouts. The Champion then splits the surface commands into two legs. The first leg commands the right aileron, right flap, throttle, and right elevator. The second leg commands the rudder, left elevator, left flap, and left aileron. The philosophy behind this surface command split is the aircraft would maintain a marginal amount of control in the event of a single component failure downstream of the Champion. Each leg is then fed into its own Acroname RX-MUX channel selector, which controls whether surface commands originate from the safety pilot receiver or the research computer. Control of the command switcher comes from the safety pilot receiver. The commands from the switcher boxes are then sent to their respective servos.

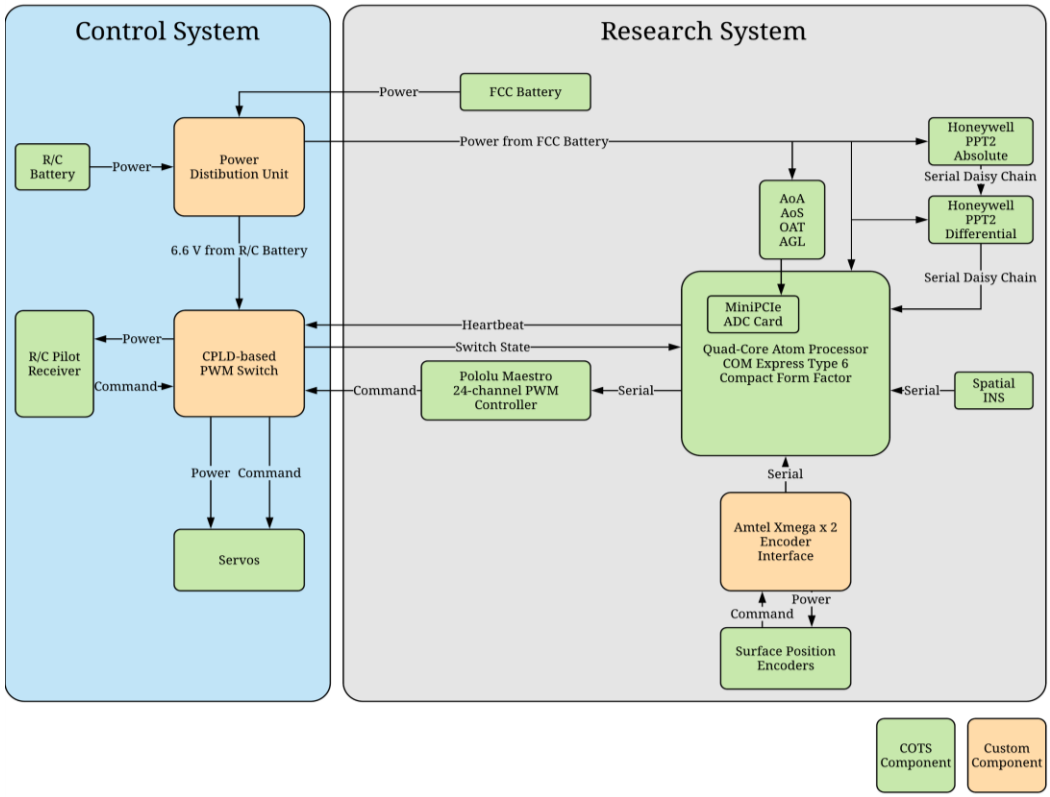

Fig. 11 Woodstock avionics schematic.

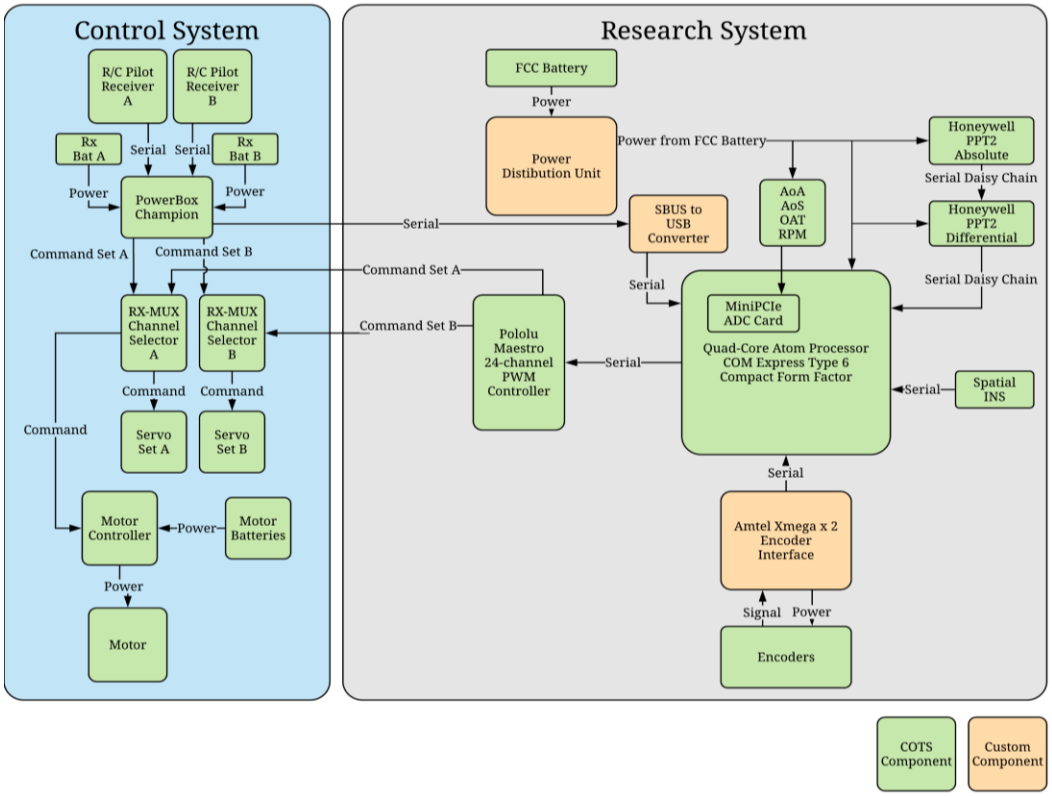

Fig. 12 E1 avionics schematic.

C. Research System

A secondary objective of the L2F project was to test the assumption that computer technology has progressed to the point where a small and inexpensive off-the-shelf system would be sophisticated enough to run complex learning algorithms in real time. The system was comprised of inexpensive sensors and an economical flight control computer (FCC) small enough to fit in the airframe but powerful enough to handle the complex learning and adaptation algorithms. The flight computer ran a Linux Ubuntu low latency operating system on an Intel Atom 1.9 GHz quad-core processor with 8 GB of RAM and a 120 GB solid state drive. Off the shelf accessories included a Pololu PWM generator and a 4-channel 16-bit A/D board. Custom components included a power distribution unit, a digital encoder reader (for control surface position), a PWM safety switch (Woodstock), and a Hall effect sensor board (for motor RPM on E1). Analog sensors included an AOA/AOS vane, outside air temperature thermocouple, a laser altimeter (Woodstock), and the RPM sensor (E1). Digital sensors were control surface position encoders, absolute and differential pressure transducers, and an inertial measurement unit (IMU). The IMU measured a 3D magnetometer, angular rates and linear accelerations and outputted both attitude and GPS solutions.

The research flight control algorithm performed at a rate of 50 Hz and recorded over 1250 variables. These variables included general flight code parameters like flags, raw encoder and analog data, and a timestamp. The algorithms ran on the flight hardware with a processor overrun margin of roughly 20%. Other variables included surface positions, sensor states, guidance and control commands, and internal parameters from the flight algorithm.

D. Avionics/Research System Integration

The primary goal of the system integration task was to ensure the research system could perform adequately without compromising the primary avionics system. The ConOps dictated that the remote pilot have ultimate control over the flight vehicle, either through direct command of control surfaces or by giving the research system permission to assume control. A requirement was that this permission could be rescinded at any time and direct control could be assumed. For the Woodstock glider, this was done by simply using a custom PWM (servo command) selector, which chose between commands from the research system or direct commands from the remote pilot. The decision was made by a dedicated R/C channel tied to a switch on the pilot's transmitter. The PWM selector switch also incorporated a watchdog timer that continuously monitored the health of the research control system. If research commands were selected and the research system dropped out for more than 750 ms, the selector switch would auto-generate flight termination commands for a left-hand spin. This could be overridden by the remote pilot at any time to assume direct control. If the R/C link was lost for more than 1 second, the R/C system would assume direct control and command a right-hand spin for flight termination. Because the remote pilot would no longer be in ultimate control during a lost-link event, termination commands would be sent even if the research system was in command and working properly. However, the remote pilot could resume command if the link was reestablished. Because of this setup, the pilot could manually engage flight termination by depowering the R/C transmitter.

The powered aircraft had a similar ConOps with the added requirement of having the research system read remote pilot commands. This was done through custom hardware that read in the Futaba SBUS (proprietary R/C command bus) and transmit it to the research system through a USB connection. It was wired in such a way so that it could not interfere with remote pilot commands going direct through the PWM switch and out to the servos. This setup allowed the research system to read in pilot commands, modify them for research objectives, and output them to the servos. If the system failed, the pilot could assume direct command at any time.

Ground testing of the system was done initially on the iron bird setup as seen in Fig. 13. Testing was eventually moved to the different airframes when they became available. Once the system was integrated and signal paths were confirmed, testing included control surface and sensor calibrations, flight termination engagement scenarios, R/C range tests, and expected response to sensor stimulus. Simulation playbacks were also performed on the actual hardware. The system would read all sensors, then overwrite the data with simulation values. The algorithm's response was recorded and compared to simulation. Algorithm cycle time was also recorded and inspected. The cycle time analysis demonstrated a margin of roughly 20% against processor overruns.

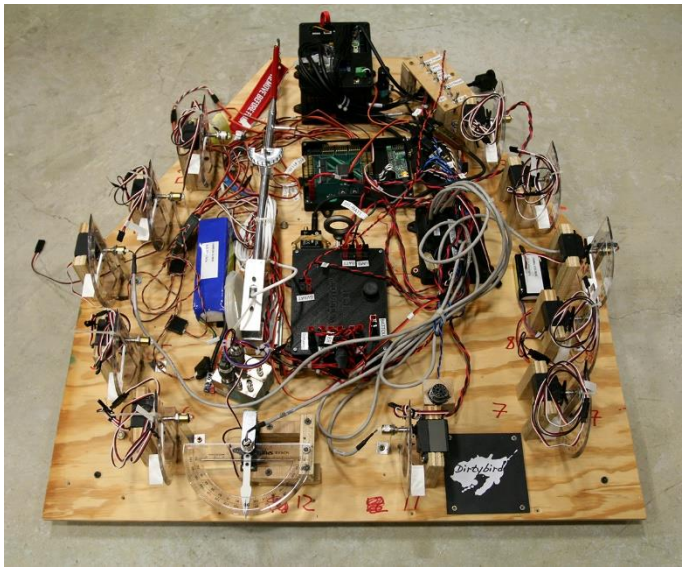

Fig. 13 Woodstock avionics on iron bird setup.

VIII. Sample Results

The data collected from flying E1 with variable stability was one of the greater successes of the project. The remote pilot could rapidly put the plane on condition, engage the stability degradation, and either fly with that degradation or let the Learn-to-Fly algorithm take full control under the constraint of the instability. Multiple attempts were made to fly each instability case by both the pilot and the L2F algorithm. Even though L2F is a learning algorithm, its memory was wiped clean in between attempts. Therefore, at the start of each attempt, the algorithm would have to maintain control of the aircraft using an initial guess on parameters, quickly adjust those parameter estimates, and improve on control performance, all while performing waypoint navigation.

An example of a test point involving the worst pitch instability flown is shown in Fig. 14, with a static margin of -16.4%. The plots start at the moment of the initialization of stability degradation. The red trace is the pilot's attempt at flying with the instability while the blue trace is an attempt by the L2F algorithm to fly that same instability. The first subplot shows the time history of pitch attitude. The pilot has difficulty maintaining level flight, which results in pitch excursions of approximately +/- 30 degrees. He also has to discontinue the stability degradation after 6 seconds because he does not want to fly through a navigational turn with the instability. The L2F algorithm initially has pitch excursions of approximately +/- 10 degrees but quickly adjusts to maintain controlled level flight before successfully negotiating two navigation turns. The initial nose-down attitude is due to a navigational correction in altitude.

The second subplot of Fig. 14 is a time history of elevator deflection. The pilot experiences control saturation in both directions while attempting to control the instability. The algorithm initially has large deflections, but they soon reduce in magnitude, remaining active to handle the unstable aircraft.

The third subplot shows the L2F algorithm's estimate of pitch stability, Cm_α. It starts with an initial guess of positive stability. It then holds that guess for one second as it observes aircraft response to control input and gathers enough data for an initial correction. After the initial observation, it quickly adjusts the estimate to one of negative stability over the course of a couple cycles and converges on a value for the change in pitching moment coefficient with respect to AOA of approximately 0.2.

The fourth subplot shows the corresponding adjustment of an effective proportional gain for the longitudinal axis, k_α. It starts at zero, and the first adjustment is in the wrong direction. This is because the initial adjustment to the pitch stability estimate is reduced but still that of positive stability. Subsequently, the gain moves in the correct direction and stabilizes once the pitch stability estimate indicates negative stability. A more comprehensive discussion on Woodstock and E1 flight results can be found in Refs. [1-4].

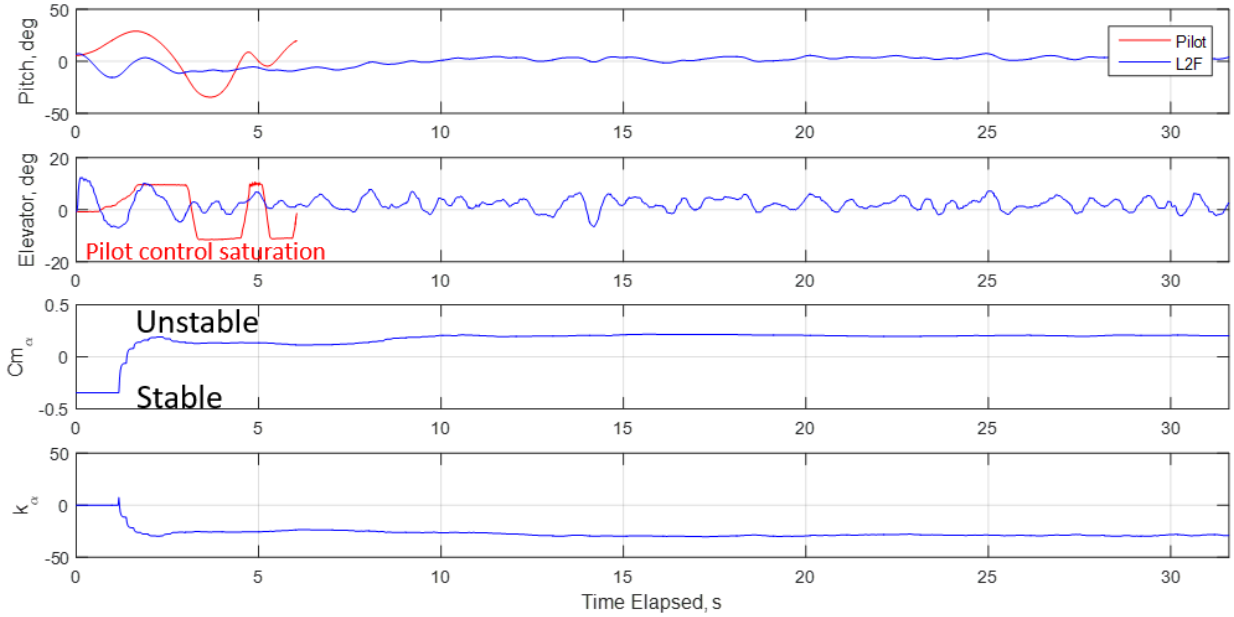

Fig. 14 Flight results with pitch instability (static margin of -16.4%).

IX. Concluding Remarks

The Learn-to-Fly test team successfully demonstrated the feasibility of learning and adapting algorithms on a series of autonomous gliders and a remotely-piloted powered aircraft. The gliders were launched via a balloon release system developed by the team. The balloon concept presented significant challenges regarding aircraft release attitude and tether avoidance. The team chose a nose-down release angle of 60 degrees while pointed away from the tether to mitigate these challenges. The team also found that a pre-marked tether line was a poor indicator of release altitude, especially with a low buoyancy margin for the system. A telemetered barometric altitude measurement was shown to be much more accurate. Despite the challenges, a balloon drop system could be efficient for reusable test articles or expendable airframes with several copies available.

The gliders themselves were flown largely in an autonomous fashion with a remote pilot serving as safety oversight. Once released, they had to perform a pull-up maneuver to achieve quasi-level flight, then perform waypoint navigation before flying a landing pattern and performing an auto-flare. The initial foam gliders showed promise, but, unfortunately, both Woodstock flights resulted in crashes. Feasibility of the L2F concept could not be assessed with either Woodstock flight. The team found that two airplane copies were insufficient when using expendable, untested airframes with highly experimental control designs.

The two Woodstock airframes were constructed using an experimental rapid prototype and design process involving 3D printed nylon. The nylon itself was shown to be too flexible and had to be reinforced. Even after reinforcement, control surfaces had to be hand-fitted to prevent binding. This lead to weight and schedule overruns. A more conventional construction process would have been quicker and more durable, in this case.

A powered aircraft was operated in a more conventional sense with a remote pilot handling takeoffs and landings while able to switch into control augmentation and autonomous modes in flight. The approach used with the powered vehicle was shown to be a very efficient method of flight testing the L2F algorithm. Eleven flights were performed on the powered airplane including a series of three instability buildup flights that was accomplished in under an hour. While the external pilot was unable to navigate with the most severe instability, the L2F algorithm rapidly recognized the stability issue, adapted its control algorithm, and successfully performed waypoint navigation.

This project as a whole was accomplished using COTS R/C avionics systems and a largely COTS research system that was comprised of inexpensive sensors and an economical FCC small enough to fit in the airframe, but powerful enough to handle the complex learning and adaptation algorithms. These algorithms could run on the flight hardware with a processor overrun margin of roughly 20%.

X. Acknowledgments

This material is based upon work funded by the NASA Aeronautics Research Mission Directorate (ARMD) Convergent Aeronautics Solutions (CAS) project under NASA's Transformative Aeronautics Concepts Program (TACP). The authors would like to thank the Learn-to-Fly test team for their hard work and dedication to this project and the Fort A.P. Hill staff for test range access and the excellent support from many individuals.

XI. References

[1] Heim, E. H., Viken, E. M., Brandon, J. M., and Croom, M. A., "NASA's Learn-to-Fly Project Overview", *AIAA Atmospheric Flight Mechanics Conference*, June 2018 (to be published).

[2] Morelli, E. A., "Practical Aspects of Real-Time Modeling for the Learn-to-Fly Concept," *AIAA Atmospheric Flight Mechanics Conference*, June 2018 (to be published).

[3] Snyder, S. M., Bacon, B. J., and Morelli, E. A., "Online Control Design for Learn-to-Fly," *AIAA Atmospheric Flight Mechanics Conference*, June 2018 (to be published).

[4] Foster, J. V., "Autonomous Guidance Algorithms for NASA Learn-to-Fly Technology Development," *AIAA Atmospheric Flight Mechanics Conference*, June 2018 (to be published).

[5] Morelli, E.A. and Klein, V., *Aircraft System Identification – Theory and Practice*, 2nd Edition, Sunflyte Enterprises, Williamsburg, VA, 2016, pp. 261-288.

[6] *The Type 'M' Kite Balloon Handbook*, Navy Department Bureau of Construction and Repair, Government Printing Office, Washington, 1919.

[7] Weinstein, R., Hubbard, J. E., and Cunningham, M. A., "Fuzzy Modeling and Parallel Distributed Compensation for Aircraft Flight Control from Simulated Flight Data," AIAA Atmospheric Flight Mechanics Conference, AIAA Aviation Forum, June, 2018 (to be published).

[8] Grauer, J. A., "A Learn-to-Fly Approach for Adaptively Tuning Flight Control Systems," *AIAA Atmospheric Flight Mechanics Conference*, June 2018 (to be published).

[9] Owens, B., Cox, D., Morelli, E., "Development of a Low-Cost Sub-Scale Aircraft for Flight Research: The FASER Project," AIAA Paper 2006-3306, June 2006. doi: 10.2514/6.2006-3306

[10] Jordan, T. and Bailey, R., "NASA Langley's AirSTAR Testbed: A Subscale Flight Test Capability for Flight Dynamics and Control System Experiments," AIAA Paper 2008-6660, August 2008. doi: 10.2514/6.2008-6660