

In-flight System Identification of the Ingenuity Mars Helicopter

Tove S. Ågren* and Allen W. Ruan†

Science and Technology Corporation, NASA Ames Research Center, Moffett Field, California, 94035

Carlos Malpica‡ and Shannah Withrow-Maser§ and Larry Meyn¶

NASA Ames Research Center, Moffett Field, California, 94035

The 68th and 69th flights of NASA’s Ingenuity Mars Helicopter marked the first dedicated system identification flights of an aerial vehicle on another planet. Chirp signals were injected into the swashplate cyclic controls for both legs of the two out-and-back flights. Frequency responses were computed from the flight data, using both the Direct Method (DM) and the Joint-Input-Output (JIO) approach, for the identification of stability and control derivatives in forward flight conditions. The resulting identified state-space models were compared against existing flight dynamics simulation models, showing excellent correlation in the higher frequency range. External disturbances were seen to introduce a bias in the identified lower frequency responses, which was partially mitigated using the JIO method. These findings will inform future modeling and flight testing efforts of Mars rotorcraft.

I. Nomenclature

A_0	=	automated sweep amplitude, deg
F, G	=	bare-airframe stability and control derivative matrices
G_{xy}	=	cross-spectra for input-to-output pair xy
\mathcal{H}	=	Hessian matrix of identification cost function
H_{xy}	=	frequency response from input-to-output pair xy
I	=	mass moment of inertia, $\text{kg} \cdot \text{m}^2$
I	=	identified parameter insensitivity, %
J	=	cost function for parametric model fit
L	=	stability and control derivatives for roll axis
M	=	bare airframe mass matrix
M_{cyc}	=	cyclic pitch control mixing matrix
m	=	mass, kg
n_ω	=	number of discretization points in frequency response
\hat{T}_c	=	composite frequency response estimate
$T_m(s)$	=	transfer function model
T_{fade}	=	time cue of linear fade-out in chirp signal, s
T_{rec}	=	length of chirp signal, s
u, v, w	=	longitudinal, lateral, and vertical velocity, m/s
V	=	wind velocity, m/s
W_g, W_p, W_γ	=	cost weighting functions for response magnitude, phase, and coherence, respectively
Y, Z	=	lateral and vertical translational stability and control derivatives
δ	=	control input signal, deg
$\hat{\gamma}_{xy}^2$	=	coherence estimator for input-to-output response pair xy
ρ	=	atmospheric density, kg/m^3
σ	=	turbulence intensity, ft/s

* Aerospace Engineer, Aeromechanics Office, tove.s.aagren@nasa.gov

† Aerospace Engineer, Aeromechanics Office, allen.w.ruan@nasa.gov

‡ Aerospace Engineer, Aeromechanics Office, carlos.a.malpica@nasa.gov

§ Aerospace Engineer, Aeromechanics Office, shannah.n.withrow@nasa.gov

¶ Aerospace Engineer, Aeromechanics Office, larry.meyn@nasa.gov

ϕ, θ, ψ	=	euler body-axis roll, pitch and yaw attitude angles, deg
$\psi_{1,eff}$	=	azimuthal location of peak cyclic magnitude, deg
θ	=	blade pitch control input, deg
$\theta_{1,eff}$	=	mean/effective cyclic input, deg
θ_{mc}, θ_{ms}	=	maximum mean cosine and maximum mean sine cyclic angle, deg
θ_h	=	heading, deg
$\omega_1, \omega_{n\omega}$	=	start and end frequency in computed frequency response, Hz
$\omega_{min}, \omega_{max}$	=	minimum and maximum frequency in chirp signal, Hz
ξ	=	kinematic coupling parameter between body-axis roll rate p and euler angle ϕ , 1/s

Subscripts

ext	=	auxiliary bare-airframe control inputs
G	=	ground frame
ls, lc	=	lower sine and lower cosine cyclic control
p, q	=	roll/pitch stability derivatives and control channels
us, uc	=	upper sine and upper cosine cyclic control
v, w	=	lateral and vertical speed stability derivatives
ϕ	=	roll angle stability derivatives

II. Introduction

THE Ingenuity Mars Helicopter became the first vehicle to demonstrate powered flight on another planet on April 19th, 2021. While designed for five test flights, the vehicle quickly surpassed initial expectations as a technology demonstrator and evolved into an aerial reconnaissance and test platform. In total, Ingenuity completed 72 flights.

The launch of Ingenuity was preceded by extensive modeling, simulation and experimental testing efforts [1–4], including a system identification (sysID) campaign to validate and calibrate a-priori physics-based flight dynamics simulation models of the vehicle. Earth-based testing introduces limitations in model identification across an entire flight envelope due to the difficulties of reproducing the true Martian flight condition experimentally. This is especially true for forward flight testing in Mars atmospheric conditions, as the capability of low-pressure wind tunnel testing is both challenging and limited. The verification and validation (V&V) campaign of Ingenuity entailed a number of tests to identify forward flight dynamics using a custom built fan array integrated in NASA’s Jet Propulsion Laboratory (JPL) 25-foot Space Simulator [5]. However, while this test suite addressed key risks in the vehicle control system design, significant modeling uncertainty in the forward flight regime remained. This uncertainty has been of particular interest to reduce since, as will be discussed in Section III.B, the stability of rotorcraft in general, and Ingenuity in particular, has a strong (adverse) dependency on horizontal airspeed. To this end, the idea of in-flight sysID was brought forth, allowing key vehicle dynamics to be characterized for the first time under true operating conditions. As Ingenuity demonstrated its robustness and performance across over its first 60 flights, sysID flight tests were considered increasingly feasible and with acceptable risk. In late fall of 2023, Ingenuity made its way up to the Martian valley Neretva Valles. The topography was ideal for a forward flight test due to its flat terrain, and the first and second extraterrestrial system identification flights were successfully executed on December 20th and 22nd, respectively.

In this paper, an overview of the vehicle properties and on-board flight control system of Ingenuity is given, including an account of the importance, and challenges, of characterizing forward flight dynamics of a Mars helicopter. Frequency-response system identification and its applicability to bare-airframe linear model identification for stability analysis and control design is discussed, highlighting the unique obstacles with closed-loop system identification testing. This is followed by an account of the design, planning, and execution of the sysID flight test campaign. Efforts in atmospheric disturbance modeling and wind estimation to support flight design and analysis are covered. Resulting frequency responses and identified linear models are presented and correlated with a-priori physics-based models. Lastly, the implications for future planetary aerial vehicle research and development are presented.

III. Flight Dynamics of Ingenuity

The flight dynamics of a Mars helicopter display significant differences compared to what would be observed on Earth. This is a direct consequence of not only the physical differences in the prevailing flight condition, but also the

mechanical design requirements of a vehicle operating in the Martian environment. A detailed account of the flight dynamics of a Mars helicopter and its influence on Ingenuity’s design is given in [1, 6]. This section gives a brief overview of the vehicle characteristics, focusing on aspects relevant to forward flight dynamics and identification thereof.

A. Vehicle Characteristics

Ingenuity features a two-bladed, counter-rotating, coaxial, hingeless rotor system. The rotor diameter measures 1.2 m, and the total mass of the vehicle is approximately 1.8 kg. Because of the low atmospheric density on Mars, approximately 1-3% of Earth’s, the tip speed is high, though limited by the low speed of sound on Mars. With the small radius, the resulting rotational speed is high, up to 2800 rpm, yielding a hover tip Mach number of $M_{tip} \approx 0.75$. Control of the vehicle is accomplished through collective and cyclic pitch control on both the upper and lower rotors, with actuator limits ranging from -4.5° to 17.5° and $\pm 10^\circ$ for the collective and cyclic pitch, respectively. Three servo motors actuate each swashplate, while two brushless, direct-drive propulsion motors spin the rotors. A more in-depth description of the vehicle’s specifications is given in [2], and a picture of the final flight vehicle is shown in Fig. 1.



Fig. 1 The Ingenuity Mars Helicopter prior to launch. Credit: NASA JPL.

The development and design of the flight control system of Ingenuity is well covered in [4, 7, 8], including the avionics architecture and implementation. What follows is a highlight of the important properties and limitations of the flight control system, comprised of the guidance, navigation, and control subsystems, in the current context of in-flight system identification test design, execution, and analysis. System identification relies on measurements of system input-output pairs of interest. For Ingenuity, this equates to time series of collective and cyclic pitch angles, and the associated bare-airframe response, respectively. Estimates of the vehicle states are provided to the control system by Ingenuity’s vision-based navigation system by filtering readings from a set of sensors including a downward-facing grayscale camera, an Inertial Measurement Unit (IMU), and a laser-range finder (LRF). The IMU consists of an accelerometer and gyroscope. In addition, pre-flight calibration of roll and pitch attitude is aided by an inclinometer. The navigation software is a velocimetry-based algorithm (MaVEN) which uses a terrain-feature mapping approach to derive state estimates. One implication of this is a strong favoring of feature-rich terrain; a lack thereof can result in degraded state estimates and even destabilization. Feature tracking also becomes increasingly difficult with increasing groundspeed. Moreover, the algorithm relies on the assumption of a planar surface model, the violation of which is also a potential source of position and heading drift. In [9], the sensitivity to terrain smoothness is discussed and pointed out as one of the limitations of the navigation system. In addition to being a source of cross-track errors, rapid terrain discontinuities tend to be falsely attributed to abrupt inertial (vertical) velocity changes. This is due to an emergent conflict in the vertical channel from LRF and camera-based estimates and inertial measurements, which will engage the control system to inject (unwarranted) feedback. Further details about the architecture and algorithms in the navigation subsystem are provided in [8].

The control system, described in [4, 7], was designed as a nested architecture with an outer translational control loop that feeds a reference to the inner attitude loop. The control design is fixed-gain across the flight envelope, with two sets of gains for two different RPM operating points depending on atmospheric density. To mitigate correlation in the control input signals, there is a static mixing of the direct swashplate to decouple the controlled input-output pairs. For instance, pitch and roll-aligned inputs, δ_q and δ_p respectively, are defined through a cyclic mixing matrix M_{cyc} and

cyclic swashplate inputs $\theta = [\theta_{lc}, \theta_{ls}, \theta_{uc}, \theta_{us}]^T$ as

$$M_{cyc} \begin{bmatrix} \delta_p \\ \delta_q \end{bmatrix} = \begin{bmatrix} \theta_{lc} \\ \theta_{ls} \\ \theta_{uc} \\ \theta_{us} \end{bmatrix} \quad (1)$$

The frequency range of interest for control is from 0.15 to 8 Hz, and is a direct consequence of the fast dynamics of the small scale vehicle. The servo actuator bandwidth is 12 Hz, with an inner-loop crossover frequency around 2.5-3 Hz.

B. Forward Flight Dynamics

A consequence of forward flight is that as the helicopter transitions from hover, the aerodynamic and inertial forces it experiences increase the coupling between lateral-longitudinal and vertical dynamics due to the azimuthal asymmetry of the rotor loads. This coupling tends to have an adverse impact on vehicle stability and flying qualities. In particular, the emergence of an angle-of-attack instability, mainly attributable to the rotor dynamics, yields an increased sensitivity to pitch moments due to vertical airspeed perturbations. Considering the low-frequency response associated with the bare-airframe dynamics, this instability expresses itself as a positive stability derivative M_w , with body-frame vertical velocity w . A detailed account of forward flight dynamics of helicopters is given in [10]. The lack of aerodynamic damping of the flapping dynamics in the thin Martian atmosphere led to the implementation of a considerable stiffening of the hingeless rotor system, including both the blades and the hub. This adaptation pushed the poorly damped flapping modes well beyond the control system's bandwidth, reducing the risk of detrimental resonant coupling which could compromise stability. A ramification of this stiffening is a large increase in hub moments from the rotor contributing to the pitch moment, resulting in an aggravated angle-of-attack instability as well as a faster transient response to gust.

As a result, Ingenuity exhibited a typical pitch-heave instability that grew rapidly with increasing advance ratio [1]. The root loci in Fig. 2 show how the characteristic longitudinal phugoid mode, marked in red, moves further into the right-half plane as the horizontal airspeed increases, as a consequence of an increasing M_w and the kinematic coupling between pitch rate and vertical acceleration. This has important implications for the control design, as it greatly affects the controllability of the system; inaccurate modeling of these dynamics may erode stability margins.




Fig. 2 Root loci of the bare-airframe dynamics with increasing horizontal airspeed.

C. Modeling and Simulation

The primary analysis tool for flight dynamics and control design in the development and analysis of Ingenuity was the Helicopter Control Analysis Tool (HeliCAT) [6], built on the DARTS/DSHELL multibody simulation framework developed at JPL [11]. The framework allows for end-to-end time domain simulations of the vehicle dynamics, featuring modeling capabilities of actuators and sensors, terrain, wind, contact dynamics, and flight software integration. Rigid body representations of the rotor, mast, landing gear, and fuselage are connected by hinges to build up the full multi-body

model. Aerodynamic forces on the rotors are computed based on blade-element theory, and a Peters-Pitt type finite state dynamic inflow model is utilized to capture global effects of the induced rotor wake on the flight dynamics. HeliCAT allows for the extraction of full and reduced order linear models around vehicle trim points used for control system design and stability analysis. The resulting linear time periodic (LTP) systems inherent to rotorcraft are reduced to linear time invariant systems through a state expansion approach using harmonic decomposition as detailed in [6].

IV. Frequency-Response System Identification

Frequency-domain methods for rotorcraft flight dynamics naturally lend themselves to stability analysis and control system design, in particular because of their applicability to unstable systems. The frequency-response method for system identification of aircraft as embodied in Comprehensive Identification from Frequency Response (CIFER) is used throughout this work, the implementation of which is covered at length in [12].

A. Methodology

At a high level, the system subject to identification is excited by input signals designed to generate a dynamic response rich in spectral content across the frequency region of interest. The collected time series are used to generate input-output frequency responses through multivariate spectral analysis, to which a parametric aircraft dynamic response model can be fit.

To determine the extent of sufficiently correlated spectral content, the coherence, $\hat{\gamma}_{xy}^2(f)$, of the frequency response of an output signal y to input x is calculated as defined in Eq. (2) where \hat{G}_{xx} , \hat{G}_{yy} , \hat{G}_{xy} , are the input-, output-, and cross-spectra, respectively. Ranging from 0 to 1, the coherence measures the fraction of the output spectrum that is linearly attributable to the input spectrum at a given frequency. Generally, a coherence greater than 0.6 indicates the response has sufficient quality over the frequency range of interest [12].

$$\hat{\gamma}_{xy}^2(f) = \frac{|\hat{G}_{xy}(f)|^2}{|\hat{G}_{xx}(f)||\hat{G}_{yy}(f)|} \quad (2)$$

A parametric model can then be fit to the associated frequency response, optimized to a coherence-weighted cost function defined by Eq. (3), with number of discrete frequency points, n_ω , over the frequency range ω_1 and ω_{n_ω} . In Eq. (3), $|T_c| - |T_m|$ and $\angle T_c - \angle T_m$ define the difference between the magnitude and phase of the data and model transfer functions. $W_g = 1.0$ and $W_p = 0.01745$ are the weights for the magnitude and phase of the frequency response, respectively. $W_\gamma(\omega) = [1.58(1 - e^{-\gamma_{xy}^2})]^2$ is a weighting function dependent on the value of the coherence function at each frequency [12]. The desired upper bound for a model fit is a cost of $J \leq 100$, and a cost $J \leq 50$ indicates a near identical fit response.

$$J = \frac{20}{n_\omega} \sum_{\omega_1}^{\omega_{n_\omega}} W_\gamma [W_g (|T_c| - |T_m|)^2 + W_p (\angle T_c - \angle T_m)^2] \quad (3)$$

The Cramer-Rao (CR) bounds of the identified model parameters define the minimum expected standard deviation of the estimate while the insensitivity, I , indicates that the parameter is not important in the selected model structure. Both are expressed as percentages of the i -th identified parameter θ_i and are defined in Eq. (4) and Eq. (5), respectively, where \mathcal{H}_{ii} represents an element in the Hessian matrix, \mathcal{H} , of the identification cost function. According to [12], $I \leq 10\%$ and $CR \leq 20\%$ reflects a sufficiently accurate identified model.

$$CR_i = \left| \frac{\sqrt{(\mathcal{H}^{-1})_{ii}}}{\theta_i} \right| \times 100\% \quad (4)$$

$$I_i = \left| \frac{1}{\sqrt{\mathcal{H}_{ii}}} \right| \times 100\% \quad (5)$$

B. Challenges of Closed-Loop System Identification

Rotorcraft are typically open-loop unstable [10]. As such, the intrinsic premise of in-flight system identification is that the control system must be engaged throughout the flight to stabilize the vehicle. Although minimal interaction from

control feedback is preferable for system identification purposes, identifying bare-airframe dynamics from closed-loop test conditions is theoretically possible [13]. This section highlights the limitations and concerns in real-life applications.

In practice, active feedback control can potentially suppress bare-airframe excitation inputs, especially in low-frequency regions below the crossover frequency of the control system. Ultimately, this will reduce the signal-to-noise ratio (SNR) and, consequently, may deteriorate the accuracy of the extracted frequency response. In addition, stabilizing feedback increases the risk of cross-correlation between control inputs, which can prohibit accurate identification of a frequency response model if not accounted for [14]. While proper control allocation and phasing to de-correlate inputs can be used to mitigate this issue, an arguably greater concern is the presence of noise and/or external disturbances, which can pose serious challenges when combined with feedback regulation. While noise to the input signal is a lesser concern in general, potential pollution of output data includes atmospheric disturbances, response to unmodeled secondary inputs or non-linearities, as well as measurement noise in sensors or errors in the navigation system outputs. In the context of Ingenuity, the navigation system's interpretation of the terrain variation as inertial state perturbations is an example of a source of measurement noise. Altogether, closing the loop introduces correlation between noise and control inputs, which in turn introduces a bias relative to the true frequency response. Fig. 3 shows a block diagram schematic of a single-input-single-output (SISO) closed-loop roll-rate response in the presence of disturbance, with mixer \mathbf{C} , feedback controller \mathbf{K} , bare-airframe plant \mathbf{P} , tracking command inputs u , bare-airframe control inputs δ_p , external bare-airframe input directly summed to the swashplate δ_{ext} , external noise n , and (measured) roll rate p .




Fig. 3 Schematics of closed-loop roll-rate p response model in the presence of noise.

The bare-airframe roll-rate frequency response, p/δ_p , can be estimated using the Direct Method (DM) through the auto and cross-spectras as

$$H_{\delta_p p} = \frac{G_{\delta_p p}}{G_{\delta_p \delta_p}}. \quad (6)$$

Under the assumption of no measurement or process noise, Eq. (6) is an unbiased estimator in the case of active feedback. Similarly, Eq. (6) is also unbiased regardless of noise level if no feedback control is engaged. However, the emergence of a non-zero auto-spectra $G_{\delta_p n} \neq 0$, due to noise being fed back through active feedback regulation, yields a biased frequency response \hat{H} . In the simplified case of unity feedback from Fig. 3, the biased frequency response $\hat{H}_{\delta_p p}$ as derived in [15] becomes

$$\hat{H}_{\delta_p p} = H_{\delta_p p} + \frac{G_{\delta_p n}}{G_{\delta_p \delta_p}}. \quad (7)$$

The second term in the right hand side of Eq. (7) is the bias error, the magnitude of which has a strong dependency on the signal-to-noise ratio. Notably, in scenarios where noise excitation predominates, the computed frequency response using the Direct Method approach may still show satisfactory coherence, potentially providing a misleading indication of the quality of the response and jeopardizing the accurate identification of the bare-airframe dynamics. In general, it is important to ensure sufficient input signal magnitude prior to the flight test to minimize the bias error. In [12], a rule-of-thumb is recommended to retain a SNR (in a Power Spectral Density (PSD) sense) above three to limit the bias error to less than 10%. This is not always possible, due to uncertainty in the test condition and/or the actuator limits. Approaches to system identification in the presence of unknown disturbances through modeling of disturbance effects on vehicle dynamics explicitly and implicitly to attain disturbance-free dynamics are presented in [16, 17]. However, when no dynamic features of the disturbance are known, such as the periodicity and magnitude, the separation of system dynamics and disturbances effects is a highly non-trivial problem; notably, the case of unknown and non-periodic disturbances is particularly ill-posed. The subsequent section presents an alternative approach to the Direct Method to resolve the issue of correlated noise bias.

C. The Joint-Input-Output method

Akaike [18] introduced the Joint-Input-Output (JIO) method as a means to address the issue of measurement noise correlation in the analysis of systems with feedback. In recent work, it has also shown great success in sysID applications with highly correlated inputs [14, 19]. The JIO method is a post-processing step to the DM. Instead of directly computing the bare-airframe response, JIO is a two-step process that includes two intermediate responses with the introduction of a reference signal. In the context of the roll-rate response and choosing the reference signal as the external swashplate input δ_{ext} as illustrated in Fig. 3, this becomes

$$H_{\delta_p P} = H_{\delta_{ext} \delta_p}^{-1} H_{\delta_{ext} P}, \text{ where} \quad (8)$$

$$H_{\delta_{ext} \delta_p} = \frac{G_{\delta_{ext} \delta_p}}{G_{\delta_{ext} \delta_{ext}}}, \quad H_{\delta_{ext} P} = \frac{G_{\delta_{ext} P}}{G_{\delta_{ext} \delta_{ext}}} \quad (9)$$

With this approach, Eq. (8) is unbiased even in the cases where $G_{\delta_p n} \neq 0$. From Eq. (9), it is noted that the overall accuracy of the JIO method is predicated upon the quality of the intermediate responses. Guidelines and considerations when using the JIO for closed-loop system identification in rotorcraft applications are detailed in [20].

V. Flight Test Design and Planning

Extensive simulation-based testing preceded the in-flight sysID campaign to support the design and planning of the flights. The ultimate objective of the simulation process was to design a flight that balanced obtaining the best data quality with sufficient risk tolerance.

A. Operational Conditions and Constraints

On-planet sysID is limited with regards to controlling the testing environment while respecting constraints on power, thermal, and actuator limits. Atmospheric density and wind conditions are uncontrolled variables, in contrast to Earth-based flight and vacuum chamber testing, and uncertainty in these variables must be taken into account. This was considered in the operational planning of the testing campaign, leading to the decision of performing "out-and-back" flights. The out-and-back trajectory provides additional information about the prevailing wind conditions, as the helicopter encounters both a potential headwind and tailwind (assuming relatively constant wind conditions throughout the flight), thus bounding the differential airspeed between the two legs. An out-and-back trajectory was also preferable from a telecommunication point of view, as the take-off and landing site are in close proximity, allowing for ease of communication with the Perseverance rover. A drawback of the two-legged flight is that the available sweep time length significantly shortens due to the additional accelerations and decelerations. Power constraints limit the maximum flight time, capped at approximately 150 seconds, resulting in a window to execute a frequency sweep in steady flight at 10 m/s of just under 30 seconds. The prevailing atmospheric density at the time window considered for the system identification flights was estimated to approximately $\rho = 0.014 \text{ kg/m}^3$ and served as the baseline value in simulations. Finally, while the full operations planning also accounted for thermal management, this does not have a direct impact on the test flight design and will not be covered herein.

B. Atmospheric Disturbance Modeling

Including the effect of noise and external disturbances on the vehicle response is integral for assessing flight risk and the extent of potential degradation in the quality of the system identification results as discussed in Section IV.B. Unsteady wind is a dominant source of disturbance, and a simple model was developed to include the effect of atmospheric turbulence on Mars in the simulation campaign. It is noted that the scope of the effort is not intended to be a highly accurate characterization of the planetary boundary layer of Mars. Rather, the primary motivation is to assess the potential implications of unsteady wind conditions on the test flight.

Conventionally, modeling atmospheric turbulence effects on aircraft flight dynamics leverages a continuous (PSD) approach. Turbulence is assumed to be frozen in time and space, isotropic and homogeneous [21]. A common choice is the Dryden model, which has spectral densities in a rational closed form. Linear forming filters, parameterized with turbulence intensity, σ , and length scale, L , are driven with Gaussian noise to generate time series of the wind velocity components, derived to replicate the theoretical PSDs [22]. The original parametrization of the Dryden model is specified for Earth applications in the Military Specification MIL-F-8785 [23] and served as a starting point for an augmented model for a Mars helicopter. The Mars 2020 rover Perseverance [24] is equipped with a suite of sensors

comprised in the Mars Environmental Dynamics Analyzer (MEDA), including both a pressure and a wind sensor [25]. While previous efforts characterized Martian vortices and dust devils using MEDA data [26, 27], this study leverages wind sensor data to fit a Dryden type model in order to generate representative wind gust time histories to inject as perturbations to the horizontal and vertical velocity channels in the full simulation environment. An upper-end estimate for turbulence intensity σ was set based on correlation with power content at low frequencies of the MEDA data. The MEDA data has a sampling frequency of up to 2 Hz, prohibiting fitting a meaningful length scale L , which was left as the default from the MIL-F-8785. Example PSDs of the longitudinal and lateral wind components in the body-frame, denoted by V_x and V_y respectively, and the time histories generated from the augmented turbulence model are presented in Fig. 4. Note that the model was not fit to this particular time-series alone but to a set of wind histories from multiple Sols. For Sol 58, it is seen to overestimate the low-frequency energy content.




Fig. 4 PSDs of MEDA wind sensor measurements of horizontal wind components from Sol 58 and wind disturbances generated from the tuned Dryden turbulence model.

C. Automated Sweep Design

The automated chirp signals (frequency sweeps) injected into the rotor cyclic inputs were designed with reference to [12]. The chirp signal has the form

$$\delta = A_0 f(\omega) \sin(\theta(t)), \quad (10)$$

with amplitude A_0 , amplitude shaping $f(\omega)$, and frequency progression $\theta(t)$ governed by

$$\theta(t) = \int^t \omega_{min} + C_2 \left(e^{\frac{C_1 \tau}{T_{rec}}} - 1 \right) (\omega_{max} - \omega_{min}) d\tau, \quad (11)$$

where ω_{min} and ω_{max} represent the sweep minimum and maximum frequency, respectively. The parameters governing the frequency progression were chosen according to conventionally recommended values $C_1 = 4.0$ and $C_2 = 0.0187$ [12]. Remaining parameter values were chosen to comply with operational constraints, including power and actuator limits as accounted for in Section V.A, while retaining acceptable coherence in the frequency region of interest. To this end, the frequency sweeps were evaluated in simulation, based on the resulting coherence in the computed frequency responses and the ability to re-construct the linear HeliCAT simulation model. Note that, conventionally, forward flight dynamics are associated with the longitudinal axis, with velocity u and pitch-rate q . However, due to a directional preference of the navigation system and the vehicle symmetry of Ingenuity, the lateral axis (with horizontal velocity v and roll-rate p) was chosen for the sysID flight test. Thus, the input-output response pair of interest for the sweep design is the on-axis roll-rate response p/δ_p .

The sweep record length, T_{rec} , is typically set by the-rule-of-thumb $T_{rec} \geq 4T_{max}$, where $T_{max} := 1/\omega_{min}$. Commonly, two full cycles at ω_{min} is included in the sweep to ensure sufficient low frequency content in the data. Theoretically, one time period of a frequency mode of interest is enough for identification purposes, but a shorter sweep length will increase the risk of response degradation caused by e.g. noise and disturbance [12]. The time window allocated for the sweep in the proposed flight sequence, limited by power constraints and acceleration/deceleration time, results in an upper bound of $T_{rec} = 22$ sec when including a padding of trimmed flight in the beginning of the sysID sequence.

Determining ω_{min} becomes a trade-off between capturing lower-frequency dynamics of interest as accounted for in Section III.B, while adhering to the allotted sweep record length T_{rec} . Given the requirements, a nominal value of $\omega_{min} = 0.1$ Hz was chosen such that $T_{rec} \approx 2T_{max}$. No full cycle inputs were included in the sweep design due to the limited sweep time and available capabilities in the flight software. The maximum frequency, ω_{max} , was set to 10 Hz in order to be beyond the inner-loop crossover frequency but below the actuator bandwidth frequency.

The frequency sweep is injected during forward flight when the trimmed cyclic control effort aligned with the direction of travel is non-zero. A large steady-state control effort increases the risk of transient control saturation, degrading the control system's ability to maintain stability. The sweep amplitude was evaluated and chosen to achieve acceptable coherence while ensuring that no sustained control saturation would occur when also considering additional cyclic efforts associated with navigation errors (e.g. due to terrain) and atmospheric disturbances. Amplitude shaping in the form of a one-period fade-in and a cued linear fade-out with respect to time around the crossover frequency was also implemented to mitigate excessive control feedback amplification and avoid sharp drops at the edges of the input signal. Figure 5 shows an example automated sweep that was injected directly into the swashplate, corresponding to $\delta_{p,ext}$ in Fig. 3, with amplitude $A_0 = 0.7$ degrees. In the presence of turbulence with intensity $\sigma_n = 10$ ft/s, Figure 5 also illustrates the actual resulting servo input when the automated sweep has been superimposed with the fed-back control signal resulting in $\delta_{p,cmd}$.




Fig. 5 Simulated designed automated sweep $\delta_{p,ext}$ and total fed-back control signal $\delta_{p,cmd}$ as commanded to the swashplate.

Without the inclusion of atmospheric disturbance in the simulation, even a small sweep amplitude of $\delta_{p,ext} = 0.4$ degrees was shown to yield coherence close to unity in the computed frequency response as illustrated in the on-axis roll rate response in Fig. 6.




Fig. 6 Simulated roll rate p response to roll aligned input δ_p with decreasing sweep amplitude A_0 .

However, as accounted for in Section IV.B, identification of the bare-airframe dynamics in closed-loop conditions must consider potential degradation due to noise and disturbances. To this end, simulations were also run with unsteady wind disturbances in HeliCAT. Note that wind time histories were not generated with the atmospheric disturbance model described in Section V.B as this model was not implemented in the HeliCAT simulation environment where the final tests were conducted. The wind disturbance model in HeliCAT does not feature any higher frequency content, as compared to the tuned Dryden-type model. As illustrated in Fig. 7, the impact of turbulence on the computed vehicle bare-airframe response is dominantly expressed as a bias and a coherence drop in the low frequency region. As the auxiliary sweep excitation $\delta_{p,ext}$ is directly superimposed with the control-gained reference trajectory, the control system interprets the sweep as a disturbance and tries to suppress it. Below the crossover frequency, the controller is successful in suppressing the "disturbance," resulting in little aircraft motion and a large noise-to-signal ratio. Above the crossover frequency, the control system is not able to suppress the aircraft motion to the sweep to the same extent, and consequently the signal-to-noise ratio increases and the frequency response bias decreases.




Fig. 7 Simulated roll rate p response to roll aligned input δ_p with sweep amplitude $A_0 = 0.6$ degrees in the presence of random wind.

To better characterize the introduced bias in the response, simulated sweeps were run with the turbulence model described in Section V.B, with an increasing turbulence intensity. The resulting frequency responses computed with the SISO Direct Method are shown in Fig. 8 and shows how the response attenuates and the coherence drops in the low frequency region as the turbulence intensity increases.




Fig. 8 Simulated roll rate p response to roll aligned input δ_p as function of modeled turbulence intensity σ_n .

To counter the degradation in the quality of the data from turbulence, the JIO method, as covered in Section IV.C, was employed to improve the extraction of the linear model from the simulated data with turbulence. Fig. 9 compares

the HeliCAT linear model with the identified models calculated from both the Direct Method and the JIO method. The magnitude of the response computed with the JIO method exhibits less attenuation in the lower frequencies compared to the Direct Method, and more closely aligns with the underlying model. Additionally, it is noted that while the Direct Method has a higher coherence below crossover, this is partially due to it being artificially inflated from the feedback control in response to the disturbances.

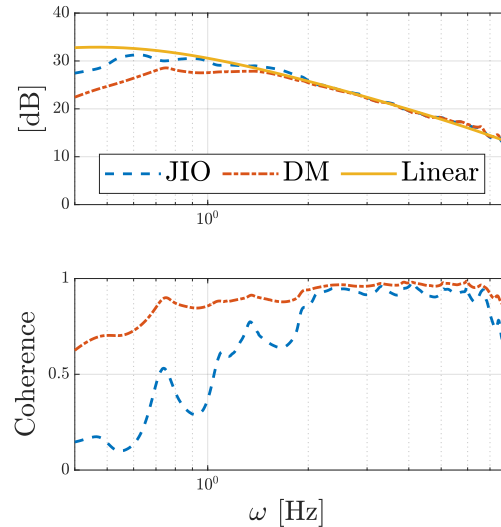


Fig. 9 Comparison of simulated frequency response with the Direct and JIO methods in the presence of high-intensity disturbance.

Having verified that the sweep would not saturate control limitations and the model extraction methodology would work, the final set of sweep parameters was determined and is presented in Table 1.

Table 1 Final sweep design parameters.

Parameter	Value
T_{rec}	22 [sec]
ω_{min}	0.1 [Hz]
ω_{max}	10 [Hz]
T_{fade}	15 [sec]
A_0	0.6 [deg]

With the parameter values in Table 1, Fig. 10 shows the required lower rotor sine cyclic control effort throughout the simulated out-and-back flight with integrated flight software, including terrain considerations and random wind disturbance. No transients exceed the actuator limit and the results do not exhibit any control saturation. The spikes observed in Fig. 10 are due to sharp changes in the commanded reference roll angle, which are induced by a commanded acceleration/deceleration.




Fig. 10 Lower sine cyclic, θ_{ls} , control effort for simulated out-and-back flight, with final sweep parameters.

D. Parametric Model Identification

The final frequency sweeps were run with the full integrated flight software in HeliCAT with parameter values as in Table 1 and commanded groundspeed of 10 m/s. A successful simulated test flight should facilitate reconstruction of the linear response and stability and control derivatives of the "true" linear model. A reduced-order, lateral-only linear state space model formulated in groundspeed coordinates was considered for derivative identification, assuming no significant cross-axis coupling. Defining state vector $x = [v_G, w_G, p, \phi]^T$, with v_G and w_G denoting lateral and vertical groundspeed respectively, and control $u = \delta_p$, the linear dynamics are modeled by

$$M\dot{x} = Fx + Gu \quad (12)$$

where the mass matrix M , stability and control derivative matrices, F and G , respectively, are defined in Eq. (13), with derivatives subject to identification highlighted in green. Note that when expressed in groundspeed coordinates, a non-zero attitude stability derivative L_ϕ emerges as a linear combination of speed derivatives, L_v and L_w , and airspeed components.

$$M = \begin{bmatrix} m & 0 & 0 & 0 \\ 0 & m & 0 & 0 \\ 0 & 0 & I_{xx} & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}, \quad F = \begin{bmatrix} Y_v & Y_w & Y_p & Y_\phi \\ Z_v & Z_w & Z_p & Z_\phi \\ L_v & L_w & L_p & L_\phi \\ 0 & 0 & \xi & 0 \end{bmatrix}, \quad G = \begin{bmatrix} Y_{\delta_p} \\ Z_{\delta_p} \\ L_{\delta_p} \\ 0 \end{bmatrix} \quad (13)$$

Without any disturbances, the identified model response shows almost perfect correlation with the true linear model, as illustrated in Fig. 11.




Fig. 11 On-axis roll rate p response of identified model from simulated sweep data without disturbance and linearized HeliCAT model.

The fit derivatives associated with the response model in Fig. 11 are presented in Table 2. While originally taken as free variables in the parametric model, the speed derivatives L_v and L_w exhibit very high insensitivities due to the lack of low-frequency content and dominance of L_ϕ in the response. Including them as derivatives to fit severely degrades the quality of the overall fit, and as a result, both were taken as model-based for the presented fits. The fit roll damping derivative L_p also exhibits Cramer Rao bound and insensitivity exceeding the guidelines, highlighted in red in Table 2. The cost $J = 13.3$ of the final fit was well below the recommended guideline $J \leq 100$.

Table 2 Identified derivatives from simulated frequency sweeps without disturbance. Density $\rho = 0.014 \text{ kg/m}^3$, cost $J = 13.3$.

	L_p	L_ϕ	L_{δ_p}
Linearized Model	0.033	0.33	8.91
Identified	0.035	0.34	8.72
CR bound, %	38.6	15.3	5.0
Insensitivity, %	19.0	5.7	1.8

Including the turbulence disturbance in the non-linear simulation results in a biased frequency response computed with DM, as illustrated in Fig. 8. A comparison of the frequency responses associated with linear models identified using the DM and JIO method, respectively, in the presence of high-intensity turbulence is presented in Fig. 12. Note that the atmospheric density was changed to $\rho = 0.0137 \text{ kg/m}^3$ for this simulation due to updated density predictions for the time of flight test. While the JIO method tends to mitigate the magnitude attenuation below crossover, the bias, when compared to the true linearized model, remains.

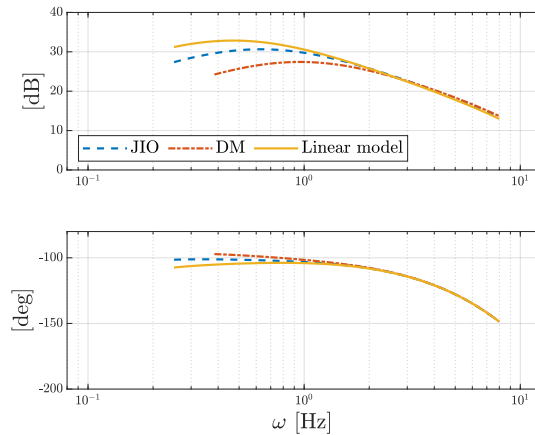


Fig. 12 Comparison of roll rate response p/δ_p of identified models from simulated sweep in presence of turbulence disturbance, using the JIO and DM method.

The associated identified derivatives are presented in Table 3. Whereas the control derivative L_{δ_p} is fit within reasonable accuracy compared to the true model, the primary effect of the introduced noise-correlation bias is a significant overprediction of the angle-of-attack derivative L_ϕ . Consistent with Fig. 12, the JIO method results in less of an overestimation and yields a more favorable fit cost. Note that roll rate damping L_p was excluded from the fit and taken as modeled due to its large insensitivity.

Table 3 Comparison of identified derivatives in the presence of disturbance using the JIO and Direct Method, $\rho = 0.0137 \text{ kg/m}^3$.

	L_ϕ	L_{δ_p}	Cost, J
Modeled	0.26	8.1	-
JIO Method	0.49	8.47	7.38
Direct Method	1.13	8.89	21.59

VI. Wind Estimation

Linearized flight dynamics models are only valid around their trim point, including environmental operating conditions such as atmospheric density and windspeed. However, since Ingenuity lacks the sensing capability to directly measure density and airspeed, both need to be indirectly inferred or measured independently. While density measurements are readily available from the MEDA data, a previous malfunction of the wind sensors on Perseverance prohibits obtaining direct wind measurements [27]. This section describes an alternative methodology to inversely estimate vehicle airspeed from the control actuator data.

An estimate of the airspeed magnitude and direction can be derived from the low frequency cyclic control effort of the vehicle [9]. By comparing the filtered in-flight cyclic and collective inputs to simulated trim cyclic and collective solutions for a variety of forward flight speeds, the approximate airspeed magnitude can be backed out through an interpolated look-up table. To correlate the cyclic inputs specifically, the root mean square of the maximum sine cyclic, θ_{ms} , and maximum cosine cyclic, θ_{mc} , inputs are calculated and referred to as the mean cyclic or effective cyclic, defined in Eq. (14).

$$\theta_{1,eff} = \sqrt{\theta_{ms}^2 + \theta_{mc}^2} \quad (14)$$

The mean cyclic input in trimmed flight at airspeeds ranging from hover to 30 m/s at an atmospheric density of $\rho = 0.014 \text{ kg/m}^3$ and rotor RPM 2700 is shown to the left in Fig. 13. The HeliCAT Ingenuity model was trimmed in increments of one m/s and the data was interpolated in between. As evident from Fig. 13, the mapping from airspeed to trim solution is non-monotonic, indicating that a single cyclic input could provide two possible airspeeds. To resolve this, evaluating the concurrent collective rotor input anchors which side of the curve the current flight condition is most likely to be in, as shown to the right in Fig. 13.




Fig. 13 Left: Mean cyclic trim solutions as a function of airspeed. Right: Mean collective trim solutions as a function of airspeed.

The wind direction can be approximated following a similar approach, making predictions based off of the rotor

cyclic inputs as derived in Eqs. (15) and (16),

$$\psi_{1,eff} = \tan^{-1} \left(\frac{\theta_{ms}}{\theta_{mc}} \right) \quad (15)$$

$$\theta_h = \psi - \psi_{1,eff} \quad (16)$$

where θ_h is the heading, ψ is the orientation of Ingenuity in the inertial frame and $\psi_{1,eff}$ is the azimuthal location of peak cyclic magnitude. Note that Ingenuity does not typically change azimuthal orientation when changing heading. Given the estimated airspeed and heading, along with the groundspeed and direction from the navigation system, the windspeed vector is determined as the groundspeed vector minus the airspeed vector.

VII. Results

Flights 68 and 69 occurred in Neretva Vallis on the Martian mornings of December 20th and 22nd, 2023, respectively. Both flights followed an out-and-back trajectory, where the first leg of both flights was in the Northeast direction, and subsequently, the return leg was in the Southwest direction as indicated in red on the map in Fig. 14. A snapshot from the Return To Earth (RTE) camera taken during Flight 68 is displayed to the right in Fig. 14.

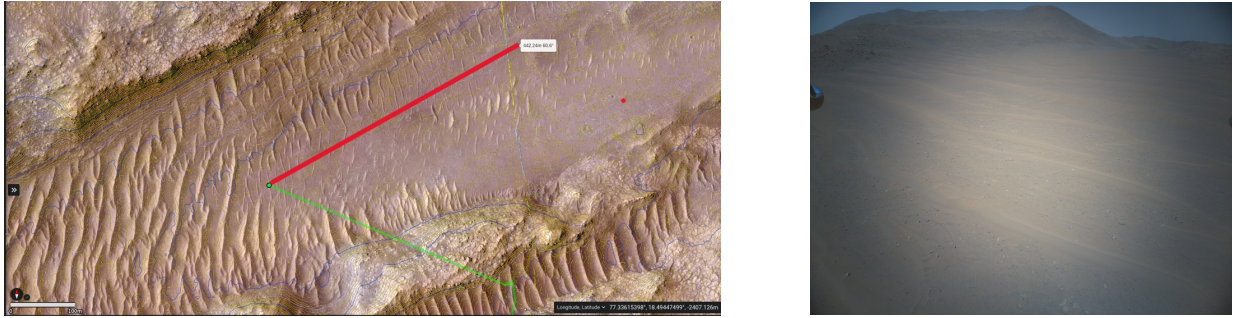


Fig. 14 Left: Top-view map of Neretva Vallis with the sysID flight trajectory marked in red. Right: An RTE camera image taken during Flight 68. Credit: NASA JPL.

Each leg included an acceleration to a constant 10 m/s lateral groundspeed before executing the designed frequency sweep excitation. Fig. 15 shows the estimated groundspeed velocity in all axes, as provided from the navigation subsystem, during the system identification sections. In addition to the 10 m/s commanded lateral groundspeed, the groundspeed heave velocity, w_G , exhibits a slight negative, constant off-set from zero, as well as an observable variation. Whether this unsteadiness is attributable to actual inertial velocity changes or noisy predictions due to variations in the terrain remains inconclusive.




Fig. 15 Groundspeed velocities estimates from the navigation system during the system identification portions of Flight 68 and 69.

As rehearsed in the simulation campaign, the sweep signal was injected into the roll-aligned inputs δ_p (predominantly the sine cyclics θ_{ls} and θ_{us}) so as to have the input excitation signal aligned with the direction of flight. The flights were nearly identical, with the exception of the sweep amplitude, where Flight 69 featured a larger amplitude of 0.7 degrees compared to 0.6 degrees in Flight 68. Time series of the unfiltered roll and pitch-aligned inputs throughout the sweep portions of the flights are shown in Fig. 16. Note that while the roll-aligned input dominated, the cross-axis pitch-aligned input was also engaged, especially in Flight 68, suggesting the presence of a notable crosswind. Fast, transient control responses in both pitch and roll channels are observed in the beginning/middle of the out portions of both flights. It is hypothesized that these spikes relate to violations of the planar assumption of the terrain.




Fig. 16 Roll and pitch-aligned control inputs during system identification portion of Flight 68 and 69.

A. Density and Wind Estimation

The average atmospheric density at the flight altitude was taken and extrapolated from the MEDA sensors on the Perseverance rover and found to be $\rho = 0.0137 \text{ kg/m}^3$ for both flights. Using the methods described in Section VI, the wind magnitudes and directions were extracted and are shown in Fig. 17. Notably, the out legs exhibit significant variability, with large magnitude and direction spikes. A drawback of the control effort based estimation approach is

that it is difficult to conclusively attribute these spikes to physical wind, since there is no way to isolate the gust response from any other noise or disturbance that has engaged the feedback control.




Fig. 17 Wind magnitude and direction for Flights 68 and 69.

To obtain a linearized model around a set trim point, a trim condition with a constant airspeed must be determined. Combining the windspeeds in Fig. 17 with the groundspeed, the airspeed in the heading direction was calculated. More specifically, the wind vector was split into components in the direction of the forward flight and in the cross-axis. It is important to note that the average value of airspeed varies depending on the approach used to compute the mean. Three different methods utilizing the windspeed estimates to obtain mean wind conditions during the sysID sequence from the washplate inputs are outlined below, and the results presented in Fig. 18 for each leg of the two flights:

- 1) Element-wise: Calculate the airspeed at every time step, and then average the airspeed across the leg at the end.
- 2) Leg-wise: Average the wind magnitude and direction for each leg (1 value for each leg), then calculate the airspeed at each timestep, and then average the airspeed across the leg at the end.
- 3) Flight-wise: Average the wind magnitude and direction for each flight (1 value for each flight), then calculate the airspeed at each timestep, and then average the airspeed across the leg at the end.

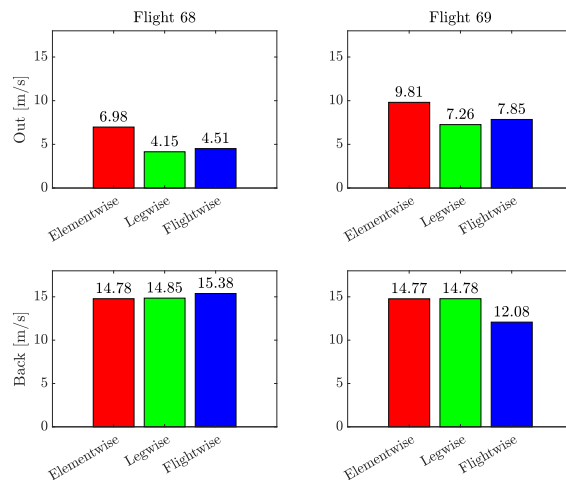


Fig. 18 Comparing methods for calculating mean airspeeds.

The approaches differ in how much the wind is assumed to change throughout the flight. The element-wise method has the smallest time resolution, and thus, heavily relies on the ability to identify wind magnitude and direction from the control inputs at shorter time scales. As a result, it has a tendency to highly weight outliers, attributing any spikes in control inputs to sharp variations in wind, which may not always be physical. Conversely, the flight-wise approach

assumes that the wind magnitude and direction remains constant for the duration of the entire flight, at the cost of not capturing potentially significant time variation in the wind condition. The leg-wise approach is a trade-off between the two, assuming the wind condition can vary from leg to leg, but not within each leg, and is the condition assumed for all further analysis. From Figs. 17 and 18, it is concluded that, on average, both out legs experienced a tailwind, while both back legs experienced a headwind, providing four separate airspeeds anchor points for model identification.

B. Frequency Responses

Unless stated otherwise, all responses shown in this section refers to a comparison of roll rate response to a roll aligned input, p/δ_p . The frequency responses from each leg for Flight 68 and 69 computed using the SISO Direct Method is presented in Fig. 19. Frequency responses were generated using CIFER’s FRESPIID with a frequency range of 0.15 Hz to 10 Hz and averaging over five separate window sizes. All four legs demonstrate a similar trend: good coherence in the higher frequency region ($\omega \geq 1$ Hz), and a loss of coherence at the lower frequencies. Out of all four, the back leg of Flight 69 demonstrates the best coherence across the entire frequency region of interest.

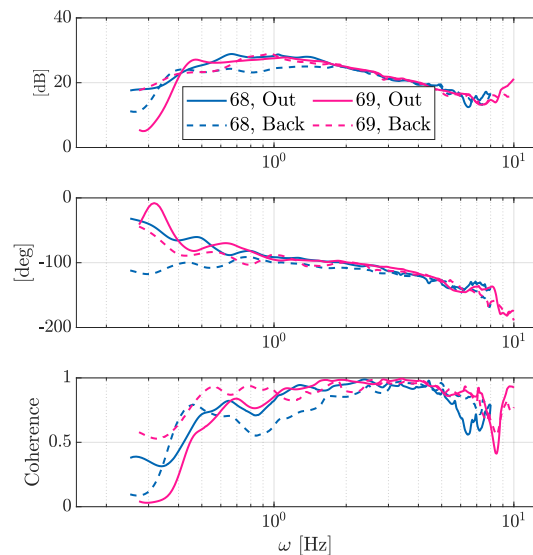


Fig. 19 Computed frequency response p/δ_p using the DM for all performed frequency sweeps.

The loss of low frequency coherence suggested potential contamination from noise and/or wind disturbances where the SNR is believed to be below the rule-of-thumb for adequate excitation. In Section V.C, the JIO method was shown to partially recover the true response in the presence of disturbance, and the same approach was applied to the flight test data and compared to the DM. Recalling the sensitivity to the quality of the intermediate responses inherent to the JIO method, this was done for Flight 69 due to the observed superior coherence and larger sweep amplitude for that particular flight. Both methods were compared to the a-priori HeliCAT models for the leg’s respective airspeed and density, and are shown in Fig. 20. It is observed that above approximately 1 Hz, the computed responses show excellent correlation with the physics-based model, in particular the magnitude, indicating that the higher-frequency response was successfully identified during the V&V campaign of Ingenuity. Below 1 Hz, the magnitude of the identified responses is seen to diverge from the model, exhibiting a more pronounced attenuation as well as phase increase. This closely resembles the observations made in Section V.C, specifically in Figs. 8 and 12, where increasing wind disturbances introduced a similar bias in the low-frequency region.

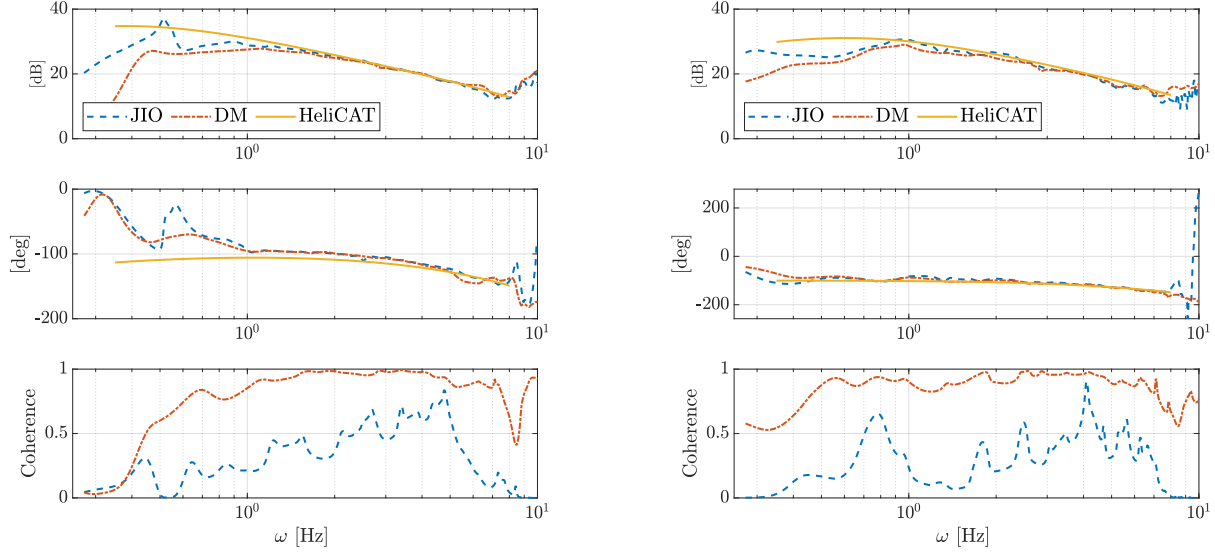


Fig. 20 Comparison of frequency response p/δ_p using the DM and JIO methods for Flight 69. Left: Out leg. Right: Back leg.

While the JIO response in Fig. 20 matches both the magnitude and phase of the linear model marginally better, the coherence is far lower than that of the Direct Method's. It is therefore hypothesized that while the coherence from the DM looks satisfactory, part of the frequency content is polluted with the presence of exogenous inputs, especially below crossover. Note, however, that the coherence estimator with the JIO method does not have a closed form expression as is the case for DM. Instead, it has to be reconstructed from the coherence of the two intermediate responses. The strategy chosen to do so for this work is to take a (weighted) minimum out of the two, which may lead to an under-prediction of the actual signal coherence [14].

C. Model Comparison

Reduced-order lateral linear models as formulated in Eq. (13) were extracted using CIFER's DERIVID state-space identification module. Due to their high insensitivities, L_p , L_v and L_w were fixed, while L_ϕ and L_{δ_p} were freed. Linear models using the DM for all legs are compared to a 10 m/s longitudinal airspeed HeliCAT model in Fig. 21. As expected, given the frequency responses the linear models are fit to, there is a consistent discrepancy between the identified models versus the HeliCAT model in the lower frequency region. Yet, the trend of a larger attenuation with airspeed (below crossover) is successfully captured in the identified models, which is expressed as a magnitude delta between the out and back flights (recalling that the back flights featured a higher horizontal airspeed due to the headwind). Theoretically, the observed bias could be an indication of a larger angle-of-attack instability than initially modeled. However, the relatively small delta, despite a substantial change in airspeed between the two flights, speaks to the fact that rather than a larger gradient in L_ϕ to airspeed than expected, the bias is more likely to be originating from disturbances and/or noise.




Fig. 21 Comparison of fit models using a SISO DM method for all frequency sweeps.

Similar to the frequency response, use of the JIO method improves the correlation of the fit parametric models with the linearized HeliCAT model in the lower frequency range. Fig. 22 shows the comparison for Flight 69, with the out leg to the left and back leg to the right. In particular, the back leg JIO model shows less attenuation in the lower frequencies, aligning closer with the physics-based model. As seen in Fig. 12, these response characteristics could be well replicated by including atmospheric disturbance in the simulation environment. This provides additional evidence that the discrepancy in observed and modeled bare-airframe roll-rate response may not conclusively be attributed to a modeling flaw, but rather that noise is being fed-back during the sweeps and is contributing to a response bias that can be partially mitigated using JIO techniques.

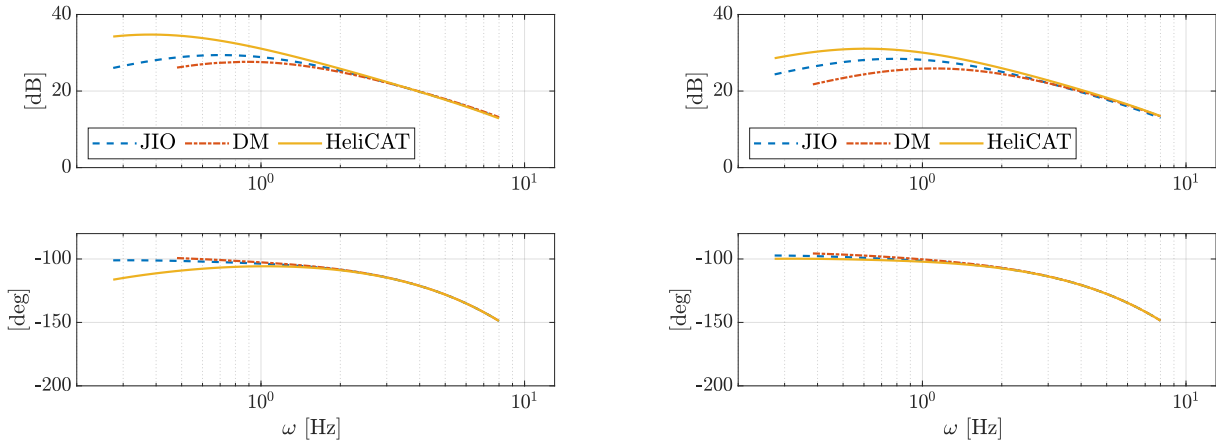


Fig. 22 Comparison of fit models using DM and JIO method for Flight 69. Left: Out leg. Right: Back leg.

D. Identification of Derivatives

The fit L_ϕ and L_{δ_p} derivatives, including cost J , insensitivity I and Cramer-Rao (CR) bounds for the out and back legs of Flight 69 are presented in Table 4 and Table 5, respectively. The data quality from Flight 68 was too low to support computing meaningful JIO responses. Identified parameters with CR bounds and/or insensitivities that exceed the rule of thumb are marked in red, which is true for the L_ϕ derivative using the JIO responses. It is also noted that the fitting cost J is consistently lower using the JIO responses.

Table 4 Fit derivatives from Flight 69, out leg.

	Modeled	DM				JIO			
	Value	J	Value	I , %	CR Bound, %	J	Value	I , %	CR Bound, %
L_ϕ	0.15	29.5	0.96	6.7	18.7	14.6	0.60	18	45.2
L_{δ_p}	8.0		8.4	2.0	5.5		8.2	3.9	9.8

Table 5 Fit derivatives from Flight 69, back leg.

	Modeled	DM				JIO			
	Value	J	Value	I , %	CR Bound, %	J	Value	I , %	CR Bound, %
L_ϕ	0.45	28.6	1.52	4.9	14.5	5.2	0.76	14.2	38.0
L_{δ_p}	8.6		8.6	2.0	6.2		8.2	4.3	11.4

Both approaches arrive at control derivatives close to the modeled value; the a-priori derivatives all fall within the Cramer-Rao bounds of the identified as shown in Fig. 23, which mirrors the excellent correlation seen in the frequency responses above cross-over in Fig. 19 and Fig. 22. The response attenuation in the lower frequencies are expressed as a significant over-estimation of L_ϕ , in particular with the DM method, which is consistent with previous observation. Whereas the values of the derivatives are off-set, both methods capture the modeled trend with airspeed, meaning that the larger velocities give a larger value for L_ϕ . This effect is overestimated in the DM results, while underestimated with JIO, when comparing to the physics-based model.

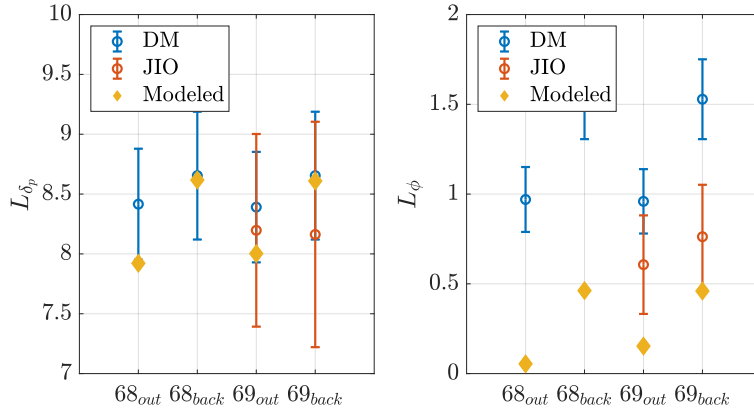


Fig. 23 Resulting fit derivatives from all frequency sweeps using both DM and the JIO method.

VIII. Conclusions

This paper detailed the motivation, design, and results from two system identification flights of the Ingenuity Mars Helicopter executed in late December of 2023. The flight tests aimed to identify key vehicle stability and control characteristics in forward flight, motivated by the lack of experimental data in this flight regime; the ultimate goal was to validate existing physics based simulation models. In addition, the effort was a first-time demonstration of a methodology for in-situ system identification of planetary aerial vehicles, setting a precedent to inform future Mars rotorcraft research efforts. The following conclusions were supported by the results:

- 1) The identified linear models from the flight data showed excellent correlation with a-priori physics based models in the higher frequency range, $\omega > 1.5$ Hz. In particular, the roll-aligned control derivative L_{δ_p} was fit within

- the Cramer-Rao bounds when compared to the linearized HeliCAT model.
- 2) Estimates of the prevailing wind conditions during the two flights were derived from cyclic control efforts. The estimates showed that the vehicle experienced a significant head (and tail) wind, reaching up to 15 m/s horizontal airspeed. No sign of instability or sustained control saturation was observed.
 - 3) While trends in stability and control derivative values with increasing airspeed were captured in the identification results, the presence of external disturbances and/or noise was seen to introduce a bias in the computed bare-airframe roll-rate frequency responses in the lower frequency region, in particular when using the Direct Method. This was primarily expressed as an overpredicted angle-of-attack derivative, L_ϕ .
 - 4) The response bias could be reduced for two of the collected time series that had sufficient coherence between the commanded reference signal and swashplate excitation when utilizing the Joint-Input-Output method, albeit with a significant reduction in input-to-output coherence in the computed bare-airframe response. The identified model using the JIO method yielded a favorable fitting cost J and a fit angle-of-attack derivative L_ϕ in better agreement with the simulation model when compared to the Direct Method.
 - 5) Notwithstanding the successful demonstration of the developed procedures for in-flight sysID, lessons for future flight testing were learned, and are listed as follows. A larger excitation signal would act to improve the SNR and improved coherence. In addition, exciting the system directly at the bare-airframe inputs, as was done in this test, likely aggravated the impact of noise; excitation through reference-tracking commands is expected to give a better SNR and a reduction in correlated noise biasing the bare-airframe response. Adding vehicle instrumentation for airspeed measurements would allow for a more accurate characterization of the vehicle trim condition and improve the accuracy of the anchoring point used for linear model extraction.

Acknowledgments

This work was carried out at NASA Ames Research Center and NASA's Jet Propulsion Laboratory (JPL), which the authors thank for their resources and support. The authors would like to express their gratitude to Martin Cacan, Joshua Anderson, and Travis Brown at NASA JPL for their invaluable efforts in planning, testing, and conducting the Ingenuity sysID flights, and in extension, to the entire Ingenuity operations and development team. Additionally, the authors would like to thank Carlos Pereyra and Hannah Dromiack for their support in tool development, Tom Berger for guidance on in-flight system identification techniques, and Wayne Johnson, Nicholas Peters, Jeremy Aires, and Sesi Kottapalli for their editorial feedback. Finally, the authors would like to acknowledge the support of William Warmbrodt and Larry Hogle during this research.

References

- [1] Grip, H. F., Johnson, W., Malpica, C., Scharf, D. P., Mandić, M., Young, L., Allan, B., Mettler, B., and San Martin, M., "Flight Dynamics of a Mars Helicopter," *43rd European Rotorcraft Forum*, 2017, p. 0014.
- [2] Balam, B., Canham, T., Duncan, C., Grip, H. F., Johnson, W., Maki, J., Quon, A., Stern, R., and Zhu, D., "Mars Helicopter Hechnology Demonstrator," *2018 AIAA Atmospheric Flight Mechanics Conference*, 2018, p. 0023. <https://doi.org/10.2514/6.2018-0023>.
- [3] Koning, W. J., Johnson, W., and Grip, H. F., "Improved Mars Helicopter Herodynamic Rotor Model for Comprehensive Analyses," *AIAA Journal*, Vol. 57, No. 9, 2019, pp. 3969–3979. <https://doi.org/10.2514/1.J058045>.
- [4] Grip, H. F., Lam, J., Bayard, D. S., Conway, D. T., Singh, G., Brockers, R., Delaune, J. H., Matthies, L. H., Malpica, C., Brown, T. L., et al., "Flight Control System for NASA's Mars Helicopter," *AIAA Scitech 2019 Forum*, 2019, p. 1289. <https://doi.org/10.2514/6.2019-1289>.
- [5] Veismann, M., Dougherty, C., Rabinovitch, J., Quon, A., and Gharib, M., "Low-density Multi-fan Wind Tunnel Design and Testing for the Ingenuity Mars Helicopter," *Experiments in Fluids*, Vol. 62, No. 9, 2021, p. 193. <https://doi.org/10.1007/s00348-021-03278-5>.
- [6] Grip, H. F., Johnson, W., Malpica, C., Scharf, D. P., Mandić, M., Young, L., Allan, B., Mettler, B., Martin, M. S., and Lam, J., "Modeling and Identification of Hover Flight Dynamics for NASA's Mars Helicopter," *Journal of Guidance, Control, and Dynamics*, Vol. 43, No. 2, 2020, pp. 179–194. <https://doi.org/10.2514/1.G004228>.
- [7] Grip, H. F., Scharf, D. P., Malpica, C., Johnson, W., Mandić, M., Singh, G., and Young, L. A., "Guidance and Control for a Mars Helicopter," *2018 AIAA Guidance, Navigation, and Control Conference*, 2018, p. 1849. <https://doi.org/10.2514/6.2018-1849>.

- [8] Bayard, D. S., Conway, D. T., Brockers, R., Delaune, J. H., Matthies, L. H., Grip, H. F., Merewether, G. B., Brown, T. L., and San Martin, A. M., “Vision-based Navigation for the NASA Mars Helicopter,” *AIAA Scitech 2019 Forum*, 2019, p. 1411. <https://doi.org/10.2514/6.2019-1411>.
- [9] Grip, H. F., Conway, D., Lam, J., Williams, N., Golombek, M. P., Brockers, R., Mischna, M., and Cacan, M. R., “Flying a Helicopter on Mars: How Ingenuity’s Flights were Planned, Executed, and Analyzed,” *2022 IEEE Aerospace Conference (AERO)*, IEEE, 2022, pp. 1–17. <https://doi.org/10.1109/AERO53065.2022.9843813>.
- [10] Johnson, W., *Rotorcraft Aeromechanics*, Vol. 36, Cambridge University Press, 2013. <https://doi.org/10.1017/CBO9781139235655>.
- [11] Lim, C. S., and Jain, A., “Dshell++: A Component Based, Reusable Space System Simulation Framework,” *2009 Third IEEE International Conference on Space Mission Challenges for Information Technology*, IEEE, 2009, pp. 229–236. <https://doi.org/10.1109/SMC-IT.2009.35>.
- [12] Tischler, M. B., and Remple, R. K., *Aircraft and Rotorcraft System Identification*, American Institute of Aeronautics and Astronautics Reston, VA, 2012. <https://doi.org/10.2514/4.868207>.
- [13] Forssell, U., and Ljung, L., “Closed-loop Identification Revisited,” *Automatica*, Vol. 35, No. 7, 1999, pp. 1215–1241. [https://doi.org/10.1016/S0005-1098\(99\)00022-9](https://doi.org/10.1016/S0005-1098(99)00022-9).
- [14] Berger, T., Tischler, M. B., Knapp, M. E., and Lopez, M. J., “Identification of Multi-Input Systems in the Presence of Highly Correlated Inputs,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 10, 2018, pp. 2247–2257. <https://doi.org/10.2514/1.G003530>.
- [15] Bendat, J. S., and Piersol, A. G., *Random Data: Analysis and Measurement Procedures*, John Wiley & Sons, 2011. <https://doi.org/10.1002/9781118032428>.
- [16] Goodzeit, N. E., and Phan, M. Q., “System and Disturbance Identification for Feedforward and Feedback Control Applications,” *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 2, 2000, pp. 260–268. <https://doi.org/10.2514/2.4548>.
- [17] Goodzeit, N. E., and Phan, M. Q., “System Identification in the Presence of Completely Unknown Periodic Disturbances,” *Journal of Guidance, Control, and Dynamics*, Vol. 23, No. 2, 2000, pp. 251–259. <https://doi.org/10.2514/2.4547>.
- [18] Akaike, H., “Some problems in the Application of the Cross-Spectral method,” *Spectral Analysis of Time Series*, 1967.
- [19] Berger, T., Tobias, E. L., Tischler, M. B., and Juhasz, O., “Advances and Modern Applications of Frequency-Domain Aircraft and Rotorcraft System Identification,” *Journal of Aircraft*, Vol. 60, No. 5, 2023, pp. 1331–1353. <https://doi.org/10.2514/1.C037275>.
- [20] Berger, T., Lopez, M. J., Wagner, A. M., and Tischler, M. B., “Guidelines for System Identification of Multicopter Vehicles with Highly Correlated Inputs,” *Vertical Flight Society 76th Annual Forum*, 2020, p. 0017. <https://doi.org/10.4050/F-0076-2020-16294>.
- [21] Durbin, P. A., and Reif, B. P., *Statistical Theory and Modeling for Turbulent Flows*, John Wiley & Sons, 2011. <https://doi.org/10.1002/9780470972076>.
- [22] Ågren, T., “Turbulence Modeling in Urban Air Mobility Applications,” Bachelor’s thesis, KTH Royal Institute of Technology, Stockholm, 2020.
- [23] Moorhouse, D. J., and Woodcock, R. J., “Background Information and User Guide for MIL-F-8785C, Military Specification - Flying Qualities of Piloted Airplanes,” *US Air Force Report AFWAL TR-81-3109*, 1981.
- [24] Farley, K. A., Williford, K. H., Stack, K. M., Bhartia, R., Chen, A., de la Torre, M., Hand, K., Goreva, Y., Herd, C. D., Hueso, R., et al., “Mars 2020 Mission Overview,” *Space Science Reviews*, Vol. 216, 2020, pp. 1–41. <https://doi.org/10.1007/s11214-020-00762-y>.
- [25] Rodriguez-Manfredi, J. A., De la Torre Juárez, M., Alonso, A., Apéstigue, V., Arruego, I., Atienza, T., Banfield, D., Boland, J., Carrera, M., Castañer, L., et al., “The Mars Environmental Dynamics Analyzer, MEDA. A Suite of Environmental Sensors for the Mars 2020 Mission,” *Space Science Reviews*, Vol. 217, 2021, pp. 1–86. <https://doi.org/10.1007/s11214-021-00816-9>.
- [26] Jackson, B., “Estimating the Heights of Martian Vortices from Mars 2020 MEDA Data,” *The Planetary Science Journal*, Vol. 3, No. 8, 2022, p. 203. <https://doi.org/10.3847/PSJ/ac87f3>.
- [27] Hueso, R., Newman, C., del Río-Gaztelurrutia, T., Munguira, A., Sánchez-Lavega, A., Toledo, D., Apéstigue, V., Arruego, I., Vicente-Retortillo, A., Martínez, G., et al., “Convective Vortices and Dust Devils Detected and Characterized by Mars 2020,” *Journal of Geophysical Research: Planets*, Vol. 128, No. 2, 2023, p. e2022JE007516. <https://doi.org/10.1029/2022JE007516>.