

# Preliminary System Identification of Dragonfly's Octocopter


Erin E. Sutton and Benjamin F. Villac

## ABSTRACT

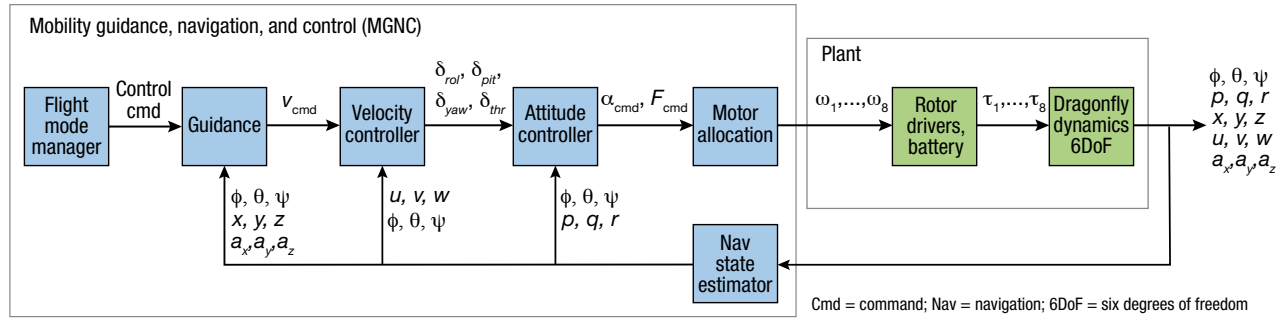
The Johns Hopkins University Applied Physics Laboratory (APL) is leading Dragonfly, a mission to study the prebiotic chemistry of Titan, one of Saturn's moons. Given Titan's diverse surface environments, mobility is crucial to the science mission, so controls engineers are faced with the challenge of designing an autonomous flight-control system for an aerial vehicle that will operate in uncertain environments. Part of the flight controller development approach involves testing with a half-scale test vehicle in an Earth environment; and one part of this process is system identification. Here, we detail the design and testing of the first round of system identification experiments with the test vehicle in which random-phase multisines were injected into the attitude commands during hover. Four experiments were performed using the half-scale test vehicle. Because of significant wind disturbances, the collected data had low coherence and were ultimately unsuitable for nonparametric frequency response estimation. System identification is an iterative process, and we present several planned ways to improve the coherence of the flight data.

## INTRODUCTION

In the mid-2030s, if all goes to plan, a Mini Cooper-sized lander will gently touch down on Saturn's moon Titan. The APL-led NASA mission is called Dragonfly, and its goals are to characterize Titan's habitability, study prebiotic chemistry, and search for signs of life.<sup>1</sup> Titan has a thick nitrogen atmosphere, methane clouds and rain, and a carbon-rich surface of lakes, rivers, craters, and dunes.<sup>2</sup> The engineering challenge is to relocate the scientific instruments to multiple high-interest sites to sample materials and measure compositions, so lander mobility is key. Taking advantage of Titan's dense atmosphere, low wind speeds, and low gravity, the lander has been designed as an uncrewed aerial vehicle (UAV), specifically an octocopter (Figure 1). It will explore terrains



**Figure 1.** Artist's rendition of Dragonfly on the surface of Titan. The scientific instruments on the self-contained rotorcraft lander will sample materials and measure composition. The lander is planned to gently touch down in the mid-2030s.



**Figure 2.** Control diagram of the lander. The mobility guidance, navigation, and control (MGNC) code is common to the Titan and Earth landers. The plant contains the Dragonfly dynamics (e.g., aerodynamics, wind conditions), the rotor drive electronics, and battery.

separated by more than 110 miles (~177,000 m) over the 3-year nominal mission duration.

But how do we design a control system for this unstable system in an uncertain environment? To start, we need a dynamics model of the UAV. Traditionally, we would build a dynamics model of the UAV and use it to develop the control system, and we would validate the model on flight-test data with the actual aircraft. That way, we could be reasonably certain that the control systems would function as expected. However, the full-size Dragonfly UAV that will be sent to Titan could not fly in Earth's atmosphere, so it is impossible to validate the model and controller response on flight-test data. Instead, we have created a dynamics model that can represent the lander on Earth or Titan with a change in model parameters. Both the model itself and the approach to generating the model parameters will be validated in the Earth environment with flight tests of a half-scale test vehicle. A comparison of the test vehicle's simulated behavior and flight-test telemetry will guide the iterative refinement of the model parameters and ultimately validate the approach to generating the model parameters.

System identification is a collection of techniques to accurately characterize the dynamic response of a physical system, subsystem, or individual component from measured data. In particular, a method called frequency domain system identification has proven useful for rotorcraft. Sinusoidal perturbations to the control inputs while the UAV is at steady-state operating conditions (e.g., hover or steady, level flight) result in changes to the UAV's attitude. By computing the frequency responses with Fourier analysis, one can identify the model parameters and uncertainties. This method has been used to analyze the handling qualities of modern military and civilian aircraft,<sup>3,4</sup> design and optimize control laws for rotorcraft,<sup>5,6</sup> estimate an aircraft's dynamic model in real time,<sup>7,8</sup> and identify the plant dynamics of the Mars helicopter.<sup>9</sup>

We have developed a process to identify the dynamics of the Dragonfly UAV using frequency system identification. It consists of (1) designing and performing the

system identification experiments, (2) preprocessing the flight data before starting system identification, (3) computing the frequency response function, (4) performing parametric modeling, (5) validating the Simulink model of the flight dynamics by comparing it with the parametric model, and (6) analyzing the sensitivity of relevant parameters. At this early stage in the project, we are iteratively improving the experimental stimuli to increase the coherence of the flight-test data (steps 1 and 2), and some of those methods and results are presented here.


## METHODS

### Simulation

The Closed Loop Sim is a Simulink (MathWorks) parameterized dynamical model of the lander. It is used for control algorithm development, sensor and actuator sizing analysis, and various design trade studies. The simulation includes parameters for Earth and Titan, and the same onboard guidance, navigation, and control (GNC) algorithms apply to both environments. It contains simplified models for all flight systems, including the plant, the system we aim to identify (Figure 2). Every experiment was thoroughly tested in the simulation prior to being loaded onto the onboard controller.



**Figure 3.** Dragonfly test UAV during flight test. The test vehicle is an approximately half-scale version of the Dragonfly lander, designed and fabricated at APL. It weighs 43.2 kg and measures 1.3 m × 1.2 m × 1.2 m, including the rotor arms, landing skids, sensors, batteries, and computing elements.



**Figure 4.** Control diagram. The stimuli were applied as perturbations to the deflection commands prior to the inner loop attitude controller.

## Test Vehicle

The test vehicle is an approximately half-scale version of the Dragonfly lander, designed and fabricated at APL (Figure 3). It weighs 43.2 kg and measures  $1.3 \text{ m} \times 1.2 \text{ m} \times 1.2 \text{ m}$ , including the rotor arms, landing skids, sensors, batteries, and computing elements.

The test article has an inertial measurement unit (HG5700, Honeywell Aerospace) running at 100 Hz, a navigation camera to measure horizontal displacements (Basler), a laser altimeter for range-to-ground sensing, and a barometric pressure sensor for altitude. It is propelled by eight rotors with electronic speed controllers, and power comes from two 5,000-mAh lithium polymer batteries. The flight software is generated directly from the closed-loop simulation's GNC code using MATLAB's Code Generation Toolbox and loaded onto the onboard processor (MCP750, Motorola).

## Experiment Design


Some of the most important decisions that need to be made during the system identification process pertain to the stimulus design. To maintain system stability, these experiments are closed loop, so the control system is actively keeping the test vehicle near the hover condition at approximately 15 m altitude. We modified the GNC code to inject periodic perturbations to the lateral, longitudinal, yaw, and collective (thrust) commands computed by the onboard controller (Figure 4). We measure the resulting Euler angles, angular velocities, and acceleration of the vehicle body frame located at the center of mass. For this analysis, we assume that the axes are decoupled, so they are treated as four independent single-input, multiple-output systems. During flight, the operator manually triggers a sequence of periodic excitations, and the axes are perturbed sequentially.

The stimulus we sequentially apply to the individual attitude commands is called a multisine

(Figure 5). The perturbation to each axis,  $\Delta\delta_{xxx}$  in Figure 4, is a linear combination of sinusoids,

$$u(t) = \sum_{k=1}^{N_f} A_k \sin(2\pi f_k t + \phi_k), \quad (1)$$

where  $N_f$  is the number of harmonics in the signal.  $A_k$ ,  $f_k$ , and  $\phi_k$  are the amplitude, frequency, and phase of  $k$ th harmonic, respectively. A good excitation signal should excite the system in the frequency band of interest. After a literature search, we decided to start with 15 frequencies: 0.1, 0.3, 0.7, 1.3, 1.9, 2.9, 3.7, 4.3, 5.3, 6.1, 7.1, 7.9, 8.9, 10.1, and 10.7 Hz. The base frequency, 0.1 Hz, is slower than the lowest bandwidth of the controller. A lower base frequency would have been preferred for studying the lower-frequency structural modes, but the experiment length is constrained by the battery life. The frequencies are prime multiples of the base frequency to eliminate harmonic coincidence and enable measurement of non-linear distortions.<sup>10</sup> The stimulus period is 10 s, so we chose an experiment length of 20 s. The integer number of periods eliminates spectral leakage, and having two periods per experiment increases the spectral resolution. In fact, more periods of input would enable time-domain averaging to reduce noise, but again, the duration was constrained by the battery capacity. Each harmonic was assigned a constant random phase between 0 and  $2\pi$ .



**Figure 5.** Input deflection command for one period of the multisine. In the time domain, the signal appears almost random. However, in the frequency domain, clear peaks are present at each of the 15 input frequencies.

The amplitude was tuned in the simulation to keep the throttle commands under 65% and limit the vertical displacement. The amplitudes were 9°, 9°, and 5° for roll, pitch, and yaw, respectively. The throttle disturbance was 20%. The axes are excited individually and sequentially, with 15 s of settling time between axes, so a single experiment lasted 265 s.

### Flight Tests

Flight tests are performed with Federal Aviation Administration approval at a private airfield. A licensed pilot closely monitors the flight and has the ability to override the flight software commands with manual joystick inputs. The guidance system ensures that the system identification commands are only accessible from hover mode. The test stimuli were preprogrammed into the flight computer and initiated by the remote operator. For this study, four experiments were performed over 2 days in July 2021.

### Preprocessing

After the flight test, we load the data into our preprocessing script in MATLAB to prepare it for the later system identification scripts. We crop the data to times during which a stimulus is applied to a given axis. We visually inspect the data to look for a response, and then we compute the magnitude-squared coherence of the angular rates and deflection commands. The coherence measures how much of the output power is linearly related to the input power. Using the roll axis as an example,

$$\gamma^2(\omega) = \frac{|S_{\phi\delta_{roll}}(j\omega)|^2}{S_{\phi\phi}(j\omega)S_{\delta_{roll}\delta_{roll}}(j\omega)}, \quad (2)$$

where  $S_{\phi\delta_{roll}}(j\omega)$  is the cross power spectral density of roll angle,  $\phi$ , and the roll deflection command,  $\delta_{roll}$ , at the input frequency,  $j\omega$ . The power spectral densities for the roll angle and roll command are  $S_{\phi\phi}(j\omega)$  and  $S_{\phi\delta_{roll}\phi\delta_{roll}}(j\omega)$ , respectively. As a rule of thumb, the minimum acceptable coherence for estimation is 0.6.<sup>3</sup> If an experiment does not reach that threshold, it is excluded from further analysis. As a final preprocessing step, the input and output data are detrended through mean subtraction and saved.

### Frequency Response Function

The broad goal of system identification is a parametric model of the underlying system, in our case,

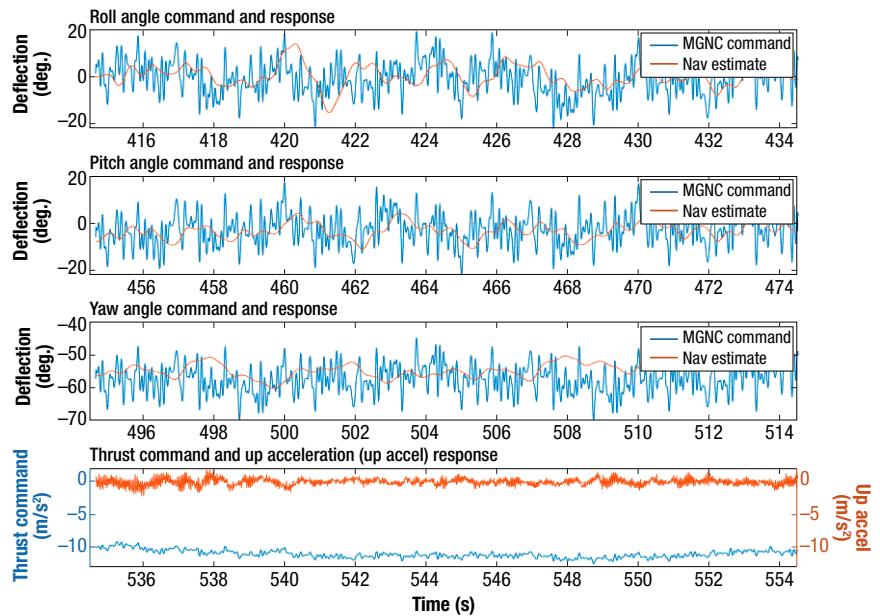
the lander. That is, we aim to find the unknown parameters  $\theta$  (e.g., mass properties, motor torque constants) in the transfer function that governs the closed-loop system dynamics,

$$G(j\omega_k, \theta) = \frac{Y(j\omega)}{U(j\omega)}, \quad (3)$$

where  $Y(j\omega)$  is the Fourier transform of the system output  $y(t)$  (e.g., roll angular rate). Similarly,  $U(j\omega)$  is the input to the system in the frequency domain (e.g., roll angle deflection command). Parametric modeling is a complicated process requiring a series of user decisions—for instance, the order of the numerator and denominator of  $G(j\omega, \theta)$ . Thus, it is recommended to obtain some knowledge about the system before performing parametric modeling. One tool that can provide information about model order, delays, and noise characteristics is the empirical, nonparametric frequency response function,

$$\hat{G}(j\omega_k) = \frac{Y(k)}{U(k)}. \quad (4)$$

Here,  $Y(k)$  is the discrete Fourier transform of the output at the  $k$ th harmonic of the input,  $U(k)$ . The frequency response function quantifies the frequency response of a system to an excitation, normalized by the magnitude of the excitation. It is essentially the transfer function of the system at discrete frequencies. Our analysis code computes the frequency response function between the angular rates and the deflection command for each axis at each stimulus frequency. In the



**Figure 6.** Raw time-domain data for experiment 3. The stimulus to the roll, pitch, and yaw commands is shown in blue. Qualitatively, the effect of the lower-frequency input stimulus is visible in the response (red). The experiment is triggered at 414 s and starts with a multi-sine disturbance to roll.



experiments presented here, the input for each axis is the deflection command, and the outputs are the corresponding Euler angle and angular rate.

## RESULTS

### Preprocessing

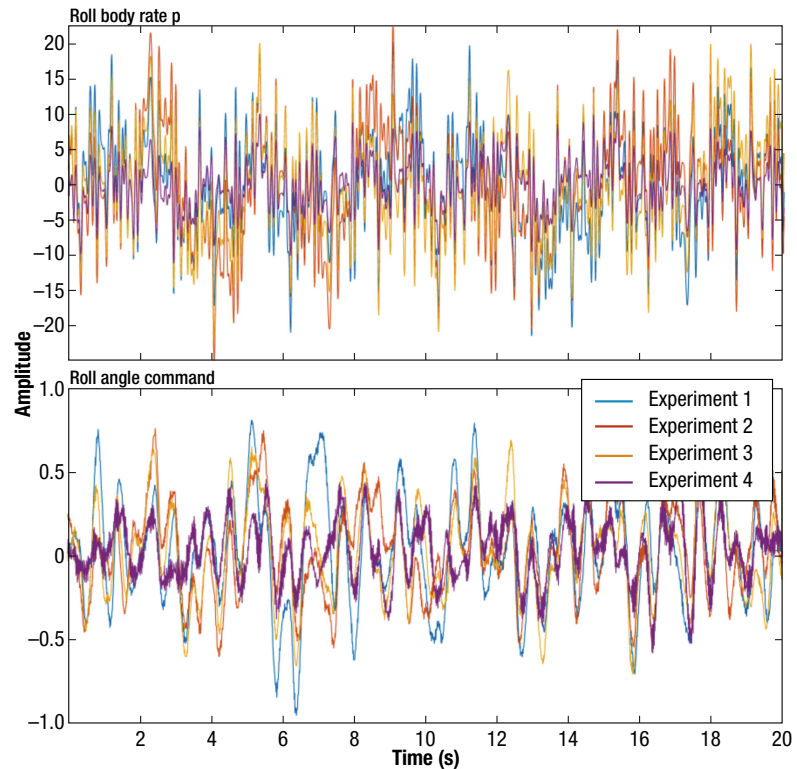
First, we visually inspected the raw data from each experiment (Figure 6). The signals appear noisy in the time domain (Figure 7 shows roll axis data). In the frequency domain, there are clear angle and angular rate responses to a few stimulus frequencies, and there is significant noise (Figure 8). Interestingly, there is variation in the deflection command between the experiments, especially at low frequencies. The roll rate amplitude response is very small compared with the input command. Again, these closed-loop experiments were conducted with the active flight-control system, so angular rates and attitude estimate from the navigation filter are fed into the inner loop attitude controller. The roll axis data were the most coherent (Figure 9), and pitch, yaw, and heave data from these four experiments were unsuitable for frequency response function estimation because of low coherence.

### Frequency Response Function


For periodic excitations, the frequency response function is obtained by dividing the output by the input spectrum (Eq. 4), and it is commonly represented as a discrete Bode plot. For this set of experiments, only the roll axis data had sufficient magnitude-squared coherence for estimation, so only the roll axis frequency response function is presented here (Figure 10).

## DISCUSSION


System identification is an iterative process, and we are investigating how to improve the coherence of the data. Literature suggests that possible causes



**Figure 7.** Time-domain data for all experiments during the roll excitation after preprocessing. There is some variation in the roll deflection command (bottom). The output roll angle (top) is not obviously coherent.



**Figure 8.** Frequency domain data for all experiments for the roll axis. The top two plots are the output data amplitude and phase, respectively. The bottom two plots are the input data amplitude and phase.




**Figure 9.** Magnitude-squared coherence at each stimulus frequency for each experiment. The minimum acceptable coherence is 0.6 (red line). Based on this result, only the roll axis was considered in further analysis, and data from experiment 4 were excluded.

**Table 1.** Mean magnitude-squared coherence over all frequencies and approximate wind speed during each experiment

	Mean( $\gamma_{p\delta}^2(\omega)$ roll)	Mean( $\gamma_{q\delta}^2(\omega)$ pitch)	Mean( $\gamma_{r\delta}^2(\omega)$ yaw)	Wind speed (m/s)
Experiment 1	0.9232	0.7491	0.5763	5.0
Experiment 2	0.9304	0.5433	0.6094	6.0
Experiment 3	0.9137	0.6045	0.524	6.0
Experiment 4	0.7687	0.6293	0.3876	6.5

Wind speed was obtained from the wunderground.com archive for a weather station near the flight-test location.



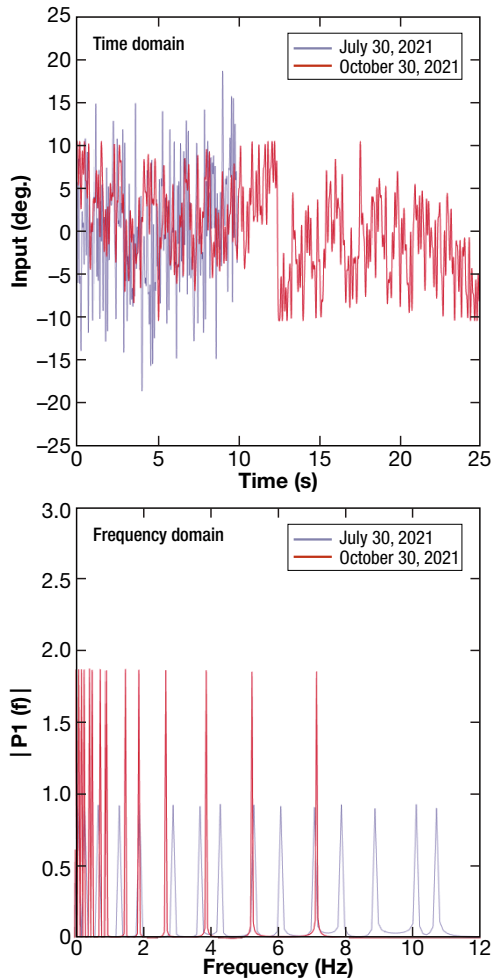
**Figure 10.** Preliminary frequency response function for the multisine stimulus showing the gain at each input stimulus frequency for the roll axis.

of low coherence are extraneous noise, leakage errors from the discrete Fourier transform (avoided with integer stimulus period), and other inputs besides  $u(t)$  (including disturbances) that contribute to the measured output.<sup>11</sup> We expect that the high wind reported on the test days explains most of the poor coherence because we also saw high experiment-to-experiment variance. Atmospheric turbulence creates transient responses and acts as an unmeasured input to the system, so tests should be conducted in calm air.<sup>4</sup> To mitigate this disturbance, future tests will be conducted on days with wind lower than 2 m/s, as measured with a portable anemometer.

The challenge in designing stimuli is exciting the system with sufficiently high power to increase the signal-to-noise ratio but not drive it from its nominal operating point. With that in mind, we made some changes to the stimuli that should improve performance in future experiments. First, we changed the frequencies in the multisine stimulus to 0.04, 0.12, 0.20, 0.28, 0.44, 0.52, 0.76, 0.92, 1.48, 1.88, 2.68, 3.88, 5.24, and 7.16 Hz. These 14 log-spaced frequencies will focus the excitations on the lower frequencies, where we expect the largest response. Next, we optimized the phases to minimize a measure called the relative peak factor.<sup>12</sup> The optimization

increased power at the input frequencies and decreased the peak-to-peak excursion of the system during an experiment, keeping the system near the hover condition (Figure 11). We have increased the battery capacity since the first round of experiments, so we will also increase the stimulus duration to 75 s and discard the first period to remove transients, which leaves two periods of data for estimation. Additionally, time-domain averaging has been shown to reduce random noise effects on the frequency response function,<sup>11</sup> so we will increase the number of experiments to 10, discard experiments with low coherence, and average the remaining experiments before applying the Fourier transform.

To maintain stability, we conducted these experiments with the control system active, so angular rates and attitude estimate from the navigation filter are fed



**Figure 11.** Comparison of the stimulus in the time domain (left) and frequency domain (right) with random-phase multisine (blue) and the proposed stimulus (red). The proposed stimulus is longer and has optimized phases, and the frequencies are log-spaced.

into the inner loop attitude controller. As a result, we observed coupling in the input commands and degraded results on single axes, especially at low frequencies (Figure 8). The coupling does not affect the frequency response function for individual channels in the frequency domain, but it does suggest that we should perform a multi-input, multi-output analysis in the future.<sup>4</sup>

## NEXT STEPS

After we complete the nonparametric frequency response analysis, the next step of the model validation process is parametric modeling. An existing analytical linearized model of the plant will be the basis for the parametric model. Then, the system identification process will fit parameters to the model that best capture the recorded flight dynamics. That will enable us to quantitatively evaluate the methods used to generate the parameter values. For instance, the value for the

rotor torque constant comes from wind tunnel experiments. If the rotor torque constant value generated during flight tests with the half-scale UAV does not match the wind tunnel data, perhaps we need to make changes to the wind tunnel test parameters (e.g., speeds, angles). In this way, experiments on Earth inform the Titan parameters, since we cannot replicate Titan in a wind tunnel. Finally, we will perform Monte Carlo simulations with small perturbations to relevant parameters and analyze sensitivities. For the most sensitive parameters, we will attempt to reduce the uncertainty with further tests. As a final step, we will build robustness into the control system around the parameters with the most uncertainty, such as those affected by unmodeled dynamics.

**ACKNOWLEDGMENTS:** We thank Rebecca Foust, Jinho Kim, and Alejandro Cabrales Hernandez for critical technical contributions. This work was carried out under NASA Contract NNN06AA01C, Task Order 80NM0018F0040.

## REFERENCES

- <sup>1</sup>J. W. Barnes, E. P. Turtle, M. G. Trainer, R. D. Lorenz, S. M. MacKenzie, et al., "Science goals and objectives for the Dragonfly Titan rotorcraft relocatable lander," *Planet. Sci. J.*, vol. 2, pp. 130–148, 2021, <https://doi.org/10.3847/PSJ/abfdcf>.
- <sup>2</sup>R. D. Lorenz, E. P. Turtle, J. W. Barnes, M. G. Trainer, D. S. Adams, et al., "Dragonfly: A rotorcraft lander concept for scientific exploration at Titan," *Johns Hopkins APL Tech. Dig.*, vol. 34, no. 3, pp. 374–387, 2018, <https://www.jhuapl.edu/Content/techdigest/pdf/V34-N03/34-03-Lorenz.pdf>.
- <sup>3</sup>M. B. Tischler and R. K. Remple, *Aircraft and Rotorcraft System Identification: Engineering Methods with Flight Test Examples*, Reston, VA: AIAA, 2012.
- <sup>4</sup>J. A. Grauer and M. J. Boucher, "Aircraft system identification from multisine inputs and frequency responses," in *Proc. AIAA SciTech Forum*, Orlando, FL, 2020, <https://doi.org/10.2514/1.G005131>.
- <sup>5</sup>S. H. Cho, S. Bhandari, F. C. Sanders, K. K. Cheung, and M. B. Tischler, "System identification and controller optimization of a coaxial quadrotor UAV in Hover," in *Proc. AIAA Scitech Forum*, San Diego, CA, 2019, <https://doi.org/10.2514/6.2019-1075>.
- <sup>6</sup>W. Wei, M. B. Tischler, and K. Cohen, "System identification and controller optimization of a quadrotor unmanned aerial vehicle in hover," *J. Amer. Helicopter Soc.*, vol. 62, 2017, <https://doi.org/10.4050/JAHS.62.042007>.
- <sup>7</sup>M. I. Alabsi and T. D. Fields, "Real-time closed-loop system identification of a quadcopter," *J. Aircr.*, vol. 56, no. 1, pp. 324–335, 2019, <https://doi.org/10.2514/1.C034219>.
- <sup>8</sup>J. A. Grauer and M. J. Boucher, "Real-time estimation of bare-airframe frequency responses from closed-loop data and multisine inputs," *J. Guid. Control Dyn.*, vol. 43, no. 2, pp. 288–298, 2020, <https://doi.org/10.2514/1.G004574>.
- <sup>9</sup>H. F. Grip, W. Johnson, C. Malpica, D. P. Scharf, M. Mandić, et al., "Modeling and identification of hover flight dynamics for NASA's Mars helicopter," *J. Guid. Control Dyn.*, vol. 43, no. 2, pp. 179–194, 2020, <https://doi.org/10.2514/1.G004228>.
- <sup>10</sup>J. Schoukens, R. Pintelon, Y. Rolain, and T. Dobrowiecki, "Frequency response function measurements in the presence of nonlinear distortions," *Automatica*, vol. 37, pp. 939–946, 2001, [https://doi.org/10.1016/S0005-1098\(01\)00037-1](https://doi.org/10.1016/S0005-1098(01)00037-1).
- <sup>11</sup>R. Pintelon and J. Schoukens, *System Identification: A Frequency Domain Approach*, Hoboken, NJ: Wiley-IEEE Press, 2012.
- <sup>12</sup>E. Morelli, "Multiple input design for real-time parameter estimation in the frequency domain," *IFAC Proc. Vol.*, vol. 36, no. 16, pp. 639–644, 2003, [https://doi.org/10.1016/S1474-6670\(17\)34833-4](https://doi.org/10.1016/S1474-6670(17)34833-4).



**Erin E. Sutton**, Air and Missile Defense Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Erin E. Sutton is a mechanical engineer at APL. She received a BS in mechanical engineering from the University of Dayton and an MS and a PhD in mechanical engineering from Johns Hopkins University. She spent a year at the Naval Air Systems Command (NAVAIR) designing flight simulators before joining APL in 2019. Her primary research interest is in enhancing existing guidance and control systems with autonomy. Her email address is [erin.sutton@jhuapl.edu](mailto:erin.sutton@jhuapl.edu).



**Benjamin F. Villac**, Space Exploration Sector, Johns Hopkins University Applied Physics Laboratory, Laurel, MD

Benjamin F. Villac is a guidance and control engineer at APL. He received a MS in mathematics and an MSE and a PhD in aerospace from the University of Michigan. After a brief postdoctoral position at Caltech and a stay at the Jet Propulsion Laboratory working on numerical methods for orbital stability analysis and low-thrust tools, he joined the faculty of the University of California, Irvine, with research focusing on astrodynamics. He subsequently joined a.i. solutions supporting the DSCOVER early operations, and he led the flight dynamics and navigation analysis team for the James Webb Space Telescope ground system. He joined APL in 2017 and has led the Dragonfly mobility subsystem team since 2018. His email address is [benjamin.villac@jhuapl.edu](mailto:benjamin.villac@jhuapl.edu).