Development and Flight Test Results of a Small UAS Distributed Flush Airdata System

ROGER J. LAURENCE III AND BRIAN M. ARGROW

Department of Aerospace Engineering Sciences, University of Colorado Boulder, Boulder, Colorado

(Manuscript received 9 November 2017, in final form 13 March 2018)

ABSTRACT

Small unmanned aircraft systems (sUAS) have proven their effectiveness for measuring both the inertial and aircraft-relative wind. One of the more common aircraft-relative wind instruments is the multihole probe (MHP). While the MHP is accurate and simple to use, two main drawbacks exist: 1) the MHP airdata system can cost several times that of the sUAS and 2) the probe itself is often exposed to damage during routine operations. Flush airdata systems (FADS) are an alternative approach and utilize pressure ports flush with the aircraft surface. This removes any external components, thereby mitigating the risk of damage to the airdata system. The work presented details the implementation of a FADS for sUAS. Computational fluid dynamics simulations were used to determine the port locations of the FADS. Airframe locations were sorted based on the total sensitivity over a range of angles of attack and sideslip. The FADS was calibrated in flight using an onboard MHP. Multilayer feedforward neural networks were employed to produce estimates of the angle of attack and sideslip, while static and stagnation ports on the fuselage measured airspeed. Results from flight testing show errors in angle of attack and sideslip were unbiased. Additionally, 97.9% of the errors in airspeed were within 1 m s^{-1} of the MHP, while 93.8% and 87.3% of the angle of attack and sideslip errors, respectively, were within 1° . Flight tests show that a FADS can be calibrated in flight and that it is an effective method for measuring the aircraft-relative wind from small UAS.

1. Introduction

When measuring the inertial wind (what an Earth-relative stationary object would experience), aircraft offer many advantages, particularly the ability to make targeted in situ observations. Volume, weight, and power requirements, as well as cost, make the integration of remote and proximal sensing systems, such as those based on radar or lidar, into small unmanned aircraft system (sUAS) airframes currently infeasible. However, sUAS are able to fly in locations where it may be unwise, impractical, or impossible for manned aircraft to fly, presenting opportunities to make in situ wind measurements that might otherwise be unobtainable. Recent sUAS missions include sampling the atmospheric boundary layer (Houston et al. 2012) and measuring the wind field in the wake of wind turbines (Kocer et al. 2011; Wildmann et al. 2014a).

The multihole probe (MHP) is a common instrument for the measurement of the aircraft-relative wind from

sUAS (Van den Kroonenberg et al. 2008; Wildmann et al. 2014b; Houston et al. 2016). A description of the transformation from the aircraft-relative wind to the inertial wind can be found in Nichols et al. (2017). Standard pitot-static probes measure only airspeed (V_∞), while the angle of attack (α ; defined as pitching up relative to the oncoming airflow) and the angle of sideslip (β ; defined as rotating right in the horizontal plane of the aircraft) are required for the full relative wind vector. Multihole probes are an adaptation of the classic pitot-static probe with additional ports surrounding the stagnation port (Telionis et al. 2009). These probes come in a variety of shapes, with differing numbers of ports, and are currently as small as 0.9 mm in diameter (Vijayagopal et al. 1998). Compressible calibration allows for a wide usable speed range for these probes (Johansen et al. 2001). The accuracy,¹ simplicity, and

¹ Errors less than 1 m s^{-1} in airspeed and 1° in angle of attack and sideslip for the Aeroprobe five-hole MHP (Aeroprobe 2016). At 20 m s^{-1} , a 1° error translates to an inertial wind magnitude error of approximately 0.35 m s^{-1} .

Corresponding author address: Roger Laurence, roger.laurenceiii@colorado.edu

FIG. 1. X-8 Skywalker with Aeroprobe five-hole MHP.

small size of the probes have led to their widespread use with sUAS. Van den Kroonenberg et al. (2008), Reuder et al. (2012), and Houston et al. (2016) report on the use of MHPs for meteorological applications. However, multiple disadvantages exist with these probes. Flow effects from the airframe can bias the measurements, and commercially available probes can be expensive (possibly several times the cost of a flight-ready sUAS). Without landing gear or a smooth runway (not all UAS field operations have access to runways, improved or otherwise), the probes are susceptible to damage during takeoff and landing. Real-world examples include clipping a wing in tall grass or skipping off a hidden rock and tumbling during landing. A multihole probe mounted on the X-8 Skywalker is shown in Fig. 1, illustrating how the mounting of the probe exposes it to potential damage during landing.

An alternative approach to measuring the relative wind is with a flush airdata system (FADS). These systems operate similarly to the MHP, in that the measured pressure from multiple ports is used to estimate the 3D relative wind vector. However, rather than employ a probe, the pressure ports of a FADS are flush with the airframe. The chief advantage of FADS for remote sUAS operations is that there is no external hardware exposed to damage, compared to an MHP. Additionally, the hardware costs of a FADS can be an order of magnitude less than the total cost of a commercial MHP.

The National Aeronautics and Space Administration (NASA) was a pioneer in the field of FADS research. Early work was done with a variety of wind tunnel models and was oftentimes concerned with subsonic or transonic speeds (Larson et al. 1980). Experimental testing of FADS was performed on a variety of aircraft shapes, sizes, and speeds, including the space shuttle (Siemers et al. 1983), an F-14 (Larson et al. 1987), an F-18 (Whitmore et al. 1990), an X-33 wind tunnel model (Whitmore et al. 1997), and a KC-135A (Larson and Siemers 1981). Multiple calibration approaches have been developed (Cobleigh et al. 1999; Haering 1995). Some of the calibration instruments, such as ground

radar or large wind tunnels, are likely inaccessible for FADS calibration on sUAS. Small MHP airdata systems, however, are often available to provide a reference for in-flight calibration.

Previous work has focused on demonstrating the efficacy of a distributed flush airdata system with wind tunnel tests (Laurence et al. 2017). The work presented in this paper differs in that flight tests were successfully performed and that the method for determining port locations was updated. Section 2 details the new method for determining the pressure port locations of the distributed flush airdata system, and section 3 presents a discussion of the hardware integration. Section 4 describes the flight tests, including an overview of the actual flights, how neural networks were used to estimate α and β , and the results from flights. This is followed by a discussion of the limitations of the system as it is currently implemented. Finally, conclusions from this project are presented in section 6.

2. CFD-based method for locating pressure ports

The X-8 Skywalker UAS is one of the primary instrument testing platforms for the Research and Engineering Center for Unmanned Vehicles (RECUV), a research center focusing on unmanned aircraft systems at the University of Colorado Boulder (RECUV 2017). The foam construction allows for easy hardware integration and quick repairs in the event of hard landings. The X-8 has a wingspan of 2.1 m and a length of 0.75 m from nose to motor mount. On a typical day (temperature around 25°C and density of approximately 1 kg m^{-3}) at 100 m above ground level (AGL) near Boulder, Colorado, with a small MHP airdata system as the payload, the X-8 cruise speed is about 17 m s^{-1} , with a corresponding endurance of about 45 min. Referring again to Fig. 1, two features of the X-8 Skywalker that are amenable for a distributed FADS are the rear-mounted pusher propeller and winglets at the wingtips. Locations closer to the leading edges of the wings are more sensitive to changes in α ; having a propeller in the rear allows the full leading edge of the wings to be used for port locations. The addition of winglets adds a portion of the wing oriented nearly perpendicular to the primary wing plane, thus increasing the sensitivity to changes in β . It is important to note that the system was designed to measure the relative wind and that while it can also be used to measure turbulence (on the scale of the airframe), this was not the objective.

Since certain locations on the airframe are relatively more sensitive to changes in flow conditions than others, it is important to determine the best locations for the FADS prior to installation. Potential flow has been used

FIG. 2. Generated mesh for CFD. Beige-colored cells are on a vertical plane section going through the midline of the model.

previously to determine the most sensitive locations on a simple hemispherical shape, such as the nose cap of an aircraft or probe (Whitmore 1991; Quindlen and Langelan 2013; Telionis et al. 2009). Analytical solutions for the pressure distribution exist but typically only for simple shapes. Computational fluid dynamics (CFD) has been used to determine port locations at a fixed spanwise distance on an airfoil (Samy et al. 2010) and across an unconventionally shaped airframe (Laurence et al. 2015).

For the present application, STAR-CCM+ (Siemens 2017) is the commercial CFD software package used for the X-8 Skywalker simulations. A polyhedral volume mesh, with 110 000 cells on the surface of the aircraft model, was generated, with cell diameters ranging from a minimum of 1.5 mm along high curvature surfaces at the leading edge to a maximum of 6.5 mm over low surface curvature parts of the airframe. Five prism cell layers are generated above the surface to help resolve the boundary layer. The generated mesh can be seen in Fig. 2. Simulations were steady state, assumed an ideal gas, and utilized the k - ϵ turbulence model (Mohammadi and Pironneau 1994, 51–62). As there are no large pressure gradients and the flow is nonseparated around the leading edge (where it is assumed the ports will be located), the k - ϵ model is an appropriate choice. Altitude (h), α , β , and V_∞ were varied, with the values shown in Table 1 (a total of 486 simulations). Freestream temperature (T_∞), pressure (P_∞), and density (ρ_∞) at both altitudes were calculated with the *U.S. Standard Atmosphere* tables (COESA 1976). For simplicity, the propeller was not included in the CFD simulations, and the simulations were used only for determining port locations.

The pressure coefficient (C_p) is utilized to incorporate the pressure response resulting from varying density (ρ_∞) and V_∞ and is defined as

$$C_p = \frac{P - P_\infty}{q_\infty}, \quad (1)$$

where q_∞ is the dynamic pressure and is defined as

TABLE 1. Simulation test matrix.

α	$[-4^\circ, -2^\circ, 0^\circ, 2^\circ, 4^\circ, 6^\circ, 8^\circ, 10^\circ, 12^\circ]$
β	$[-6^\circ, -4.5^\circ, -3^\circ, -1.5^\circ, 0^\circ, 1.5^\circ, 3^\circ, 4.5^\circ, 6^\circ]$
V_∞	$[14, 17, 20] \text{ m s}^{-1}$
h	$[1830, 2440] \text{ m}$

$$q_\infty = \frac{1}{2} \rho_\infty V_\infty^2. \quad (2)$$

The X-8 Skywalker operates at low enough speeds ($M \approx 0.05$) to be in the incompressible range. Because of this, it is assumed that C_p is unaffected by changes in speed or density. The values of C_p for each ρ_∞ and V_∞ (for a specific cell with fixed α and β) are averaged together into a mean pressure coefficient variable. All subsequent results are presented in terms of the mean pressure coefficient.²

Point 67914 (shown in Fig. 4) will serve as an example to demonstrate the approach to port selection. The response of C_p to α and β is shown in Fig. 3a. As the mean pressure coefficient has been reduced to a function of α and β , the gradient of C_p is composed of an α component, $(\nabla C_p)_\alpha$, and a β component, $(\nabla C_p)_\beta$. The gradient of C_p shows how sensitive the location is to changes in the flow angles. For point 67914, $(\nabla C_p)_\alpha$ is shown in Fig. 3b, while $(\nabla C_p)_\beta$ is shown in Fig. 3c. It is immediately apparent that the pressure coefficient is an order of magnitude more sensitive to changes in α than it is to changes in β . It is also important to note that $(\nabla C_p)_\alpha$ goes to zero when $\alpha \approx 5^\circ$. Around this angle of attack value,³ the pressure coefficient is extremely insensitive to changes in α , and errors in the pressure measurements will lead to larger errors in the estimation of α . As such, locations that experience a maximum/minimum in the pressure coefficient are excluded from consideration.

Integrating the magnitude of both $(\nabla C_p)_\alpha$ and $(\nabla C_p)_\beta$ leads to a scalar “total sensitivity” value for both the α and β directions:

$$\begin{aligned} \alpha \text{ total sensitivity} &= \int_{\alpha} \int_{\beta} |(\nabla C_p)_\alpha| d\beta d\alpha \\ \beta \text{ total sensitivity} &= \int_{\alpha} \int_{\beta} |(\nabla C_p)_\beta| d\beta d\alpha \end{aligned} \quad (3)$$

²For simplicity, the mean pressure coefficient will be represented by C_p throughout the rest of the paper. When the standard pressure coefficient is used, the altitude and airspeed will also be specified.

³Term $(\nabla C_p)_\beta$ also goes to zero but at $\alpha \approx 3^\circ$.

FIG. 3. Point 67914 CFD results. (a) Mean C_p , (b) the α component of the mean C_p gradient, and (c) the β component of the mean C_p gradient.

This is equivalent to computing the volume under the $(\nabla C_p)_\alpha$ and $(\nabla C_p)_\beta$ surfaces (i.e., the volume under Figs. 3b and 3c). Locations with a higher sensitivity magnitude over a larger area in (α, β) space will lead to larger total sensitivity values. The individual CFD mesh points on the surface of the aircraft are ranked based on their total sensitivity values. The α total sensitivity is shown in Fig. 4a, while the β total sensitivity is shown in Fig. 4b. The maps of the total sensitivity show that the leading edges of the wings are most sensitive to changes in α , while the leading edges of the winglets are most sensitive to changes in β . It is also interesting to note that the nose of the X-8 Skywalker is not the most sensitive region. This raises the possibility that the performance of FADS may be improved if the ports are located on the wings as opposed to the nose of the aircraft, as is commonly done.

In addition to locations that experience maxima/minima in the pressure coefficient, the entire underside of the aircraft and the joints where the wing transitions to the winglet are excluded. The X-8 Skywalker does not have landing gear; as such, any ports on the underside are prone to damage upon landing. The winglet joints are excluded because of the possibility of flow irregularities. With the exclusion zones in place, the mesh points are sorted by their α and β total sensitivities. Eight locations were chosen based on the α total sensitivity, with another eight locations chosen with the β sensitivity. All 16 final locations are displayed in Fig. 5. From the figure, some asymmetry in the port locations is noticeable: this is likely the result of asymmetry in the original computer-aided design (CAD) model or small errors in the simulation results.

The dynamic pressure, defined in Eq. (2), is equal to the difference between the stagnation (or total) pressure P_t and the freestream pressure P_∞ . By measuring P_t and P_∞ , the airspeed can be calculated. From the definition of C_p , it is known that at a location where the pressure equals the stagnation pressure, $C_p = 1$; for locations

equal to freestream static pressure, $C_p = 0$. The pressure coefficient at the expected trim condition ($[2^\circ, 0^\circ]$; the angles of the aircraft during steady, level flight) was used to determine the static and stagnation port locations. From CFD results, the location on the airframe where the pressure equals stagnation pressure is located on the nose, as shown in Fig. 6a, while there are multiple areas where the pressure equals the freestream static pressure, shown in Fig. 6b. When there are multiple suitable locations, the location requiring the shortest amount of tubing is usually the most appropriate (Haering 1995). The installed static and stagnation ports are displayed in Fig. 7b.

3. Assembly

The pressure sensor chosen for the FADS is the MS8607-02BA01 sensor (TE Connectivity 2015). The MS8607 sensor was chosen for its ability to measure pressure, temperature, and humidity (PTH); its low cost ($\sim \$10$ per sensor); and its high resolution for a micro-electromechanical system (MEMS)-style sensor. This is also an absolute pressure sensor; while a differential pressure sensor can result in greater accuracy, differential sensors are significantly larger than the MS8607. Space is at a premium within the wings of the Skywalker, and a smaller sensor package was given higher priority. Additionally, differential pressure sensors would require running pneumatic tubing from the wings to the fuselage, as the measured pressures are referenced against the static pressure. With the current setup, only electrical wires are passed from the wings to the fuselage, allowing the connections to be simpler. Relevant technical specifications for the MS8607 can be found in Table 2.

The circuit boards were designed to house four MS8607 sensors each. Two layers of acrylic seal the sensors onto the circuit board, separating the MS8607 sensors into individual chambers. Small ports in the top

FIG. 4. Total sensitivity for (a) α and (b) β along with location of point 67914.

layer of acrylic are used to house brass tubing; this brass tubing connects to flexible urethane tubing, which is run to the pressure port locations on the surface of the aircraft.

Each circuit board is tested upon completion of the sealing process. To verify the pressure seals, the brass tubing is sealed, and the entire circuit board is cooled to $\sim 5^\circ\text{C}$. If the sensors are sealed properly, the pressure inside will decrease upon cooling (and will increase as the circuit board is allowed to return to room temperature). The pressure from each sensor is recorded upon removal from the refrigerator, and any sensors that measure the ambient pressure (instead of increasing pressure) are known to be sealed improperly. Additionally, a rapid opening of the seal on the tubing should yield an immediate change to ambient pressure; this serves as another check of the integrity of the sensor seals. After circuit board testing was completed, it was determined that one circuit board had all four sensors sealed properly, while four other boards each had three sensors sealed properly. The board with four sealed sensors is housed in the fuselage and is dedicated to the static and stagnation ports; the remaining boards are dedicated to the ports on the wings and winglets.⁴

⁴Even though 16 port locations were chosen, there are only 12 usable sensors for those ports.

FIG. 5. Port locations from CFD results. (a) Ports 1–4 (right wing) and (b) ports 5–8 (left wing) were chosen with α total sensitivity, while (c) ports 9–12 (right winglet) and (d) ports 13–16 (left winglet) were chosen from β total sensitivity.

Small holes were drilled in the fuselage at the static and stagnation port locations. Flexible urethane tubing with an inner diameter of 3/64 in. (1.2 mm) was then installed through the drill holes. Smaller diameter ports allow for finer resolution measurements, and the port size has minimal impact on the accuracy (Paćes et al. 2010). Port sizes as small as 0.76 mm have been used with FADS in the wind tunnel (Larson et al. 1980). Unfortunately, the port size relative to the body does have an impact on the maximum incidence angle at which the error in the total pressure measurement is within 1% (Gracey 1958). From Gracey (1958), it is estimated that the FADS total pressure measurement will be within 1% of the truth for incidence angles less than 10° . The port sizing and tubing length play an important role in the distortion of pressure signals (Whitmore and Leondes 1991); however, because of the slow sampling rates, this distortion is ignored. Future work should investigate the importance of port sizing for this application, along with potential errors in the static pressure measurements.

Hardware installation is shown in Fig. 7. Tubing for the static and stagnation ports is displayed in Fig. 7b. The stagnation port is in the center of the nose, while the static ports are on the side of the nose. Space constraints necessitated the installation of the pressure boards close to the thicker portion of the wing near the leading edge, as seen in Fig. 7a. All excess tubing was cut to be flush with the aircraft surface. The pressure board and tubing are all installed on the underside of the wing for simplicity and to maintain aerodynamic efficiency. Individual tubing channels are carved into the foam wing to bury the tubes to minimize the effect on the aerodynamics of the wing. Channels for the tubing are carved into the outside surface of the winglets and are shown in Fig. 7c. Installation is finished with the application of tape over the tubing and pressure boards, in order to maintain a smooth surface.

FIG. 6. Mean pressure coefficient from Skywalker CFD simulations. The color bar has been adjusted to highlight (a) areas where pressure is equal to stagnation pressure (red) and (b) areas where pressure is equal to freestream static pressure (green).

4. Results

a. Flight tests

The Skywalker UAS, outfitted with the distributed FADS and an Aeroprobe MHP (Aeroprobe 2016) as the reference instrument, was flown on 22 March 2017. Flight testing occurred at the Boulder Model Airport in Boulder. Two flights were performed, and each lasted nearly 40 min. At the beginning of flight 1, various maneuvers were performed for validation of the FADS. This validation portion lasted 500 s and yielded 3143 data points. The remainder of flight 1 (1700 s and 10 710 data points) was used for calibration of the FADS. The entirety of flight 2 (2160 s and 13 552 data points) was also used for calibration of the FADS.

Neural networks, the regression technique used for estimating α and β from the pressure measurements, typically do not extrapolate well outside of the training domain (Wythoff 1993). To increase the training domain size, a variety of flight conditions were explored during the calibration flights. The first flight was predominately

box-shaped patterns that varied the airspeed and altitude. The majority of the second flight was circular orbits, with varying airspeed, altitude, and turning direction. In the end, approximately 60% of the calibration flight time was under manual control, with the remaining 40% under autopilot control. Because of the area where flights took place, the Skywalker was restricted to an area approximately $400\text{ m} \times 600\text{ m}$, which limited the length of any straight legs.

The validation portion of the Skywalker flight was intended to mimic real-world flight profiles. Validation was made up of transects (in the form of rectangles because of space constraints), ascending and descending orbits, and straight legs with varying altitude. Nine distinct patterns are found in the validation flight data: spiral 1, spiral 2, ellipse 1, box 1, box 2, box 3, box 4, leg 1, and leg 2. In addition to the overall validation results, results from individual patterns will also be analyzed in an attempt to determine ideal flight patterns for minimizing the errors in the FADS.

The values of α and β measured by the reference MHP during calibration are exhibited in Fig. 8a. The contour lines

FIG. 7. (a) Hardware installation on the wings along with port locations. Lower wing shows hardware installed on the wing underside. (b) Stagnation and static port locations. (c) Winglet port locations and tubing routing.

TABLE 2. Technical specifications of the MS8607-02BA01 PTH sensor.

Operating range	10 to 2000 hPa, 0% to 100% RH, −40° to 85°C
High-resolution mode	1.6 Pa, 0.04% RH, 0.01°C
Absolute accuracy (25°C)	±200 Pa, ±3% RH, ±1°C
Relative accuracy (25°C)	10 Pa
Response time	<5-ms pressure, 5-s RH

represent the number density of recorded data points. The grid spacing for the number density is 0.2° in the α direction by 0.3° in the β direction. The most common air-relative orientation during calibration is around $\alpha = 3^\circ$ and $\beta = 1^\circ$. Additionally, 80% of the flow angles are within the outermost contour line, which approximately covers the area of 2° to 6° for α and -2° to 6° for β . With the calibration data points as such, it is hypothesized that the FADS will be most accurate within this range. The recorded angles during the validation portion are shown in Fig. 8b, with the calibration data contours overlaid. The validation data show a similar distribution to the calibration data, but they are shifted toward slightly lower angles of attack.

b. Neural network training

Neural networks have been successfully implemented with FADS previously, including flight data from an F-18 (Rohloff 1998), a small UAS (Quindlen and Langelaan 2013), a mini air vehicle (Samy et al. 2010), and a wind tunnel model (Crowther and Lamont 2000). Hornik et al. (1989) have shown that a single hidden layer in a multilayer feedforward network (a class of neural networks) is enough for the network to be a universal approximator, given sufficient hidden units and training data. Thus, neural networks are known to be an excellent choice for use with flush air-data systems.

There are multiple ways to affect the accuracy of the network, with the two more obvious ways being the training method used and the hidden layer size. Different training methods use different minimization approaches to determining the weights/biases of the network by minimizing the error between the network output and the target data.⁵ The three training methods investigated are the scaled conjugate gradient method (Moller 1993; Charalambous 1992), which reduces the training time and is better suited for larger problems thanks to lower memory requirements; Levenberg–Marquardt (Hagan

and Menhaj 1994; Marquardt 1963), which employs a nonlinear least squares algorithm and is very efficient for smaller networks (up to a few hundred weights⁶); and Bayesian regularization (Foresee and Hagan 1997; MacKay 1992), which adds a smoothing term to improve generalizability⁷ but at the expense of increased computation time and memory usage. The hidden layer size represents the number of hidden units (also referred to as “neurons”) in the hidden layer. This is analogous to the degree of a polynomial fit. Too few hidden units, and the network will not appropriately capture the complexity between the inputs and outputs; too many hidden units, and the network may overfit the data.

There are several other ways to affect the accuracy of the network. The networks were initialized with the Nguyen–Widrow algorithm (Nguyen and Widrow 1990; MathWorks 2017a), which utilizes randomness in choosing the initial weights and biases. During network training, the calibration data are split into different categories:⁸ training data, which are used for adjusting the weights and biases; validation data, which are used to measure the network response to novel inputs and determine when to stop training; and testing data, which have no effect on the network and are only used to provide independent data to assess network performance. As the calibration data are typically randomly split into the different categories, reinitializing the network will change the initial weights/biases and how the data are split into the different categories, both of which can lead to different network results. Additionally, what inputs are used heavily influence the accuracy of the network. The networks trained for this project use the differential pressure from all 12 sensors (sensor pressure minus freestream static pressure), density,⁹ free-stream static pressure, and the indicated airspeed as inputs.

Ensembles of 500 networks were trained using the three different training methods and hidden layer sizes from 2 to 12 hidden units. The root-mean-square error (RMSE) was used to compare the performances of the different networks:

⁶ The number of weights for a single hidden layer feedforward network is equal to (hidden layer size $\times$ number of inputs) + hidden layer size + (number of outputs $\times$ hidden layer size) + number of outputs.

⁷ A network generalizes well when it performs as well with novel inputs as it does with training inputs (Foresee and Hagan 1997).

⁸ Not all training methods split the data into all three categories.

⁹ Density was computed with the ideal gas law, without any humidity effects.

⁵ For this project the target data are the angles of attack and sideslip measured by the reference MHP.

FIG. 8. MHP α and β data density during (a) calibration and (b) validation flights. Density contours are identical in (a) and (b), and are based on the calibration data.

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{Y}_i - Y_i)^2}, \quad (4)$$

where $\hat{Y}_i$ is the network output (or mean output from the network ensemble) and Y_i is the measured value from the reference MHP. The networks are trained with the calibration data, but the RMSE values are computed with the data from the validation portion of the flights. Levenberg–Marquardt and Bayesian regularization have comparable RMSE values, but the scaled conjugate gradient is less accurate in every case. Of the different hidden layer sizes tested, nine hidden units resulted in the greatest performance overall. While performance was comparable, the focus on generalizability from Bayesian regularization was more important than the shorter training times of Levenberg–Marquardt.¹⁰ Two thousand new networks with nine hidden units were trained with Bayesian regularization, and the single network with the lowest RMSE value was chosen. This network has RMSE values for the calibration data of 0.557° and 0.795° for α and β , respectively, and RMSE values for the validation data of 0.651° and 0.871° for α and β , respectively. The layout of this network is presented in Fig. 9.

Regression plots showing how well the network fits the calibration data are shown in Fig. 10. The horizontal axes show the target values (measurements from the reference MHP), and the vertical axes show the network

outputs. For a network that perfectly reproduces the target data, all the data would lie on the black line. The red lines are the actual linear fits between the network outputs and the target data. The linear fits appear to more closely align for the β outputs than with the α outputs. However, RMSE_α is lower than RMSE_β . This discrepancy is explained by the wider variability in the β outputs. The network outputs for β are centered about a closer to perfect line (slope closer to 1, lower mean), but they are spaced farther apart from that line (higher RMSE, higher variance).

c. Flight results

While neural networks are used to produce estimates of the angle of attack and sideslip, the airspeed is calculated from the dynamic pressure (difference between the stagnation and static pressure) in Eq. (2). Rearranging leads to the analytic solution for airspeed:

$$V_\infty = \sqrt{\frac{2(P_t - P_\infty)}{\rho}}. \quad (5)$$

Indicated airspeed (IAS) uses the standard mean sea level density, ρ_0 , in Eq. (5), while true airspeed uses the true value of ρ . All mentions of airspeed refer to indicated airspeed, unless otherwise specified. Additionally, individual errors in α , β , and IAS (ε_α , ε_β , and ε_{IAS}) are simply

$$\varepsilon_i = \hat{Y}_i - Y_i \quad (6)$$

with the notation the same as before. A positive error indicates the FADS estimate is higher than the measurement from the MHP.

¹⁰ On a 2.4-GHz, 32-core computer, training 500 networks with nine hidden units took 20 min for Levenberg–Marquardt and 70 min for Bayesian regularization.

FIG. 9. Network schematic. The transfer functions are hyperbolic tangent sigmoid for tansig (MathWorks 2017c) and linear for purelin (MathWorks 2017b).

The regression plots and RMSE values only describe how the network performs overall. Knowing how the network performs under different flight conditions (e.g., low α , positive instead of negative β , etc.) can inform the user of flight patterns that minimize the network error. Maps of the RMSE values in α - β space using the calibration data are shown in Fig. 11, with RMSE_α in Fig. 11a and RMSE_β in Fig. 11b. The calibration data are grouped into bins 0.3° in the α direction by 0.75° in the β direction. The data points that fall into each bin are then used to calculate RMSE_α and RMSE_β . Focusing on just the network performance, it is clear that lower errors occur in the general $\alpha = 2^\circ$ to 4° and $\beta = 0^\circ$ to 2° area. This is the same vicinity as the densest location of calibration data points. Overall, the network performs better with respect to α in the $\alpha = 2^\circ$ to 5° and $\beta = -2^\circ$ to 8° region, while $\alpha = 1^\circ$ to 5° and $\beta = -1^\circ$ to 4° are better with respect to β performance. It is also interesting to note that RMSE_α is sensitive to the true α but that it is relatively insensitive to the true β ; the inverse case is true for RMSE_β , where the errors are sensitive to the true β , but they are relatively insensitive to the true α .

In addition to the 2D error maps, it is helpful to see how network errors are affected by the flight parameters. Errors in IAS, α , and β are plotted against the true α in Fig. 12, along with linear fits to the errors. Clearly, the errors in the α estimate from the FADS are influenced by the true angle of attack. This tendency was also apparent in the RMSE_α map. Interestingly, the true α has virtually no effect on the errors in β from the FADS. The slope of the fits (how sensitive a FADS error is to a particular flight parameter) is presented in Table 3. For example, a 1° increase in the true α will lead to a change of 0.1962° in the magnitude of the FADS angle of attack error, ε_α , but only a change of 0.0067° in the angle of sideslip error ε_β . However, a single degree change in the true β will change the magnitude of ε_β by 0.0697° . This shows that ε_β is 10 times more sensitive to changes in the true β than to changes in the true α . Table 3 lists the sensitivities of ε_α , ε_β , and ε_{IAS} to changes in true values for IAS, α , β , ρ , and P_s .

Overall, the Skywalker's flush airdata system performed well compared to the MHP standard, as evidenced by the histogram of errors in Fig. 13. While there

FIG. 10. Comparison of target data (MHP) and output data (FADS estimate) for the neural network.

FIG. 11. (a) $RMSE_{\alpha}$ and (b) $RMSE_{\beta}$ of FADS estimates during calibration flights.

is a slight bias in the indicated airspeed measurements from the FADS, both α and β are unbiased. The bounds for the goals of 1 m s^{-1} in IAS and 1° in α and β are overlaid in Fig. 13. For IAS, 97.9% of the errors from the FADS are within the set goal; for α and β , the amounts are 89.0% and 77.5%, respectively. Unsurprisingly, the variance in the β errors is larger than the variance in the α errors (this goes back to the previous point of β having a better regression fit but worse RMSE).

FIG. 12. Effects of true α on FADS errors with the validation flight data.

The results presented so far have focused on the full validation dataset. As mentioned previously, the validation portion of the flight is composed of nine distinct patterns. Looking at the patterns individually will inform the user as to what patterns are best to fly for the system as calibrated. The best pattern overall was the first box pattern flown, displayed in Fig. 14. This is a box pattern with left-hand turns, a nearly constant altitude, lasted 55 s, and contains 345 measurements. It is immediately apparent that the FADS (orange) tracks the MHP values (blue) very well for IAS, α , and β . As has been the case so far, the FADS does a better job estimating α than β , as evidenced by the RMSE values of 0.567° and 0.693° , respectively. The measurements in β are less noisy compared to α , so it is easier to see how well the FADS β estimate tracks the truth. The largest errors in β occur with the peaks at high β values. It is already known that the FADS has larger errors for high β values. The worst overall pattern was the ellipse, where the aircraft was banking for a large portion of the pattern. The RMSE values for all of the validation patterns can be found in Table 4, along with RMSE values for the full validation dataset.

5. Discussion and system limitations

The results show the flush airdata system performs comparably to the reference multihole probe, yet the FADS, as implemented, still has some limitations. The largest limitation currently is access to a high-quality reference instrument for truth data that can be used with sUAS during calibration. The MHP has its own measurement errors inherent to the system. Additionally, flow effects from the Skywalker airframe will affect the MHP measurements. While accounting for these errors will lead to improvements in the accuracy of the FADS,

TABLE 3. FADS error sensitivity.

	IAS (m s^{-1})	α (deg)	β (deg)	ρ (kg m^{-3})	P_s (Pa)
ε_α (deg)	-0.0495	-0.1962	0.0364	-11.445	-1.56×10^{-4}
ε_β (deg)	0.0876	0.0067	-0.0697	11.531	8.485×10^{-5}
ε_{IAS} (m s^{-1})	0.0881	0.0131	-0.0573	19.297	7.783×10^{-5}

modeling these errors is beyond the scope of this paper. Plus, some of the error corrections for the MHP can be made postflight, which requires only that the network to be retrained, as opposed to requiring new flight data.

Several other limitations of the FADS, as implemented, exist. The MHP records data at 50 Hz, while the FADS records only at 6.2 Hz. The MHP is able to measure features on a shorter time scale than the FADS is able to respond to, leading to errors. This is a limitation of the current data logging setup, and different logging hardware will lead to sampling improvements. The FADS also measures the wind on a different spatial scale compared to the MHP; the MHP is on the order of millimeters, while the FADS is on the order of a meter. While neural networks have proved their capability, they offer no robustness to sensor failure, as sensor inputs cannot be excluded from the network. One possibility is training new networks without inputs from

certain sensors, but this becomes infeasible quickly. The number of new networks required to account for n_{faulty} sensors out of n_{total} sensors is

$$\text{net}_{\text{required}} = \frac{n_{\text{total}}!}{(n_{\text{total}} - n_{\text{faulty}})!n_{\text{faulty}}!}. \quad (7)$$

Accounting for a single sensor failure requires 12 new networks, while two sensor failures require an additional 66 networks. This lack of robustness to sensor failure has been mitigated by grouping sensors together and ignoring entire groups if one sensor is faulty (Rohloff et al. 1999), preprocessing the inputs with an autoassociative network (Crowther and Lamont 2000), or taking a least squares approach (Whitmore et al. 1992).

One final limitation of note is related to the airframe of the Skywalker UAS. The aircraft is made of foam, and the spar in the wings does not connect to the winglets.

FIG. 13. Histogram of errors from the validation flight data.

FIG. 14. Detailed summary of the box 1 validation pattern.

This led to noticeable flexing at the winglet joint compared to the rest of the wing. The flexing likely contributes to the larger errors in β compared to errors in α .

6. Conclusions

A new approach to designing a flush airdata system for small UAS was presented. Computational fluid dynamics was used to determine the most sensitive locations on the airframe, which led to the locations of the pressure ports. From the CFD results, it was shown that

horizontal surfaces are more sensitive to changes in α , while vertical surfaces are more sensitive to changes in β . It was also shown that the nose is not necessarily the most sensitive location on the airframe and that a flush airdata system might be more appropriate with the pressure ports located elsewhere on the airframe.

The FADS on the X-8 Skywalker was successfully calibrated in flight. An Aeroprobe MHP served as the reference instrument against which the system was calibrated. Airspeed measurements were made from measuring the dynamic pressure, while angle of attack and sideslip estimates were estimated using feedforward neural networks. Referring back to the FADS goal, 97.9%, 93.8%, and 87.3% of the airspeed, angle of attack, and angle of sideslip estimates, respectively, from the FADS were within 1 m s^{-1} and 1° of the MHP. Using the full validation flight data, the errors in α and β were unbiased. Viewing the errors in α - β space showed that the system was most accurate in the region with the highest density of training points. Therefore, flights should be designed so that the aircraft operates in the higher-accuracy area most often; alternatively, if a different higher-accuracy area is desired, then the system can be recalibrated.

The flush airdata system presented in this paper is only the first iteration of the system on the X-8 Skywalker.

TABLE 4. RMSE of the validation patterns.

	IAS (m s^{-1})	α (deg)	β (deg)
Spiral 1	0.5229	0.6143	0.7633
Spiral 2	0.3678	0.6281	0.9663
Ellipse 1	0.5019	0.7415	0.9596
Box 1	0.3400	0.5675	0.6932
Box 2	0.4314	0.6775	0.9270
Box 3	0.4359	0.6323	0.8637
Box 4	0.4720	0.6753	0.9172
Leg 1	0.3904	0.5631	0.8018
Leg 2	0.4456	0.6965	0.7714
Overall	0.4230	0.6507	0.8710

Several limitations of the system, as implemented, were discussed. Some of these limitations, such as the slow sampling rate and no humidity measurements, will be addressed in the future by installing different hardware; other limitations can be addressed through software changes. Flight results have shown the system is well suited for wind sensing applications, and improvements to the system will increase its usability for other meteorological applications.

Acknowledgments. A large number of people contributed to the success of this project. The authors acknowledge the help of Abhilash Manjunath for designing the code to collect the MS8607 sensor data; Thomas Lillis for implementing the logging software; Dan Hesselius for piloting the Skywalker and for providing general hardware support; Steve Borenstein for the circuit board design and sensor support; Chris Choate, Chris Brown, and Will Finamore for aircraft support; and the rest of the RECUV/IRISS support team. The authors are also grateful for the support of the NSF National Robotics Initiative (NRI; Award IIS-1527919; PI: Eric Frew) and BlackSwift Technologies.

REFERENCES

- Aeroprobe, 2016: Aeroprobe Micro Air Data System user manual. 2nd ed. Revision D, Aeroprobe Corporation, 38 pp.
- Charalambous, C., 1992: Conjugate gradient algorithm for efficient training of artificial neural networks. *IEEE Proc.*, **139G**, 301–310, <https://doi.org/10.1049/ip-g-2.1992.0050>.
- Cobleigh, B., S. Whitmore, E. Haering Jr., J. Borrer, and V. Roback, 1999: Flush airdata sensing (FADS) system calibration procedures and results for blunt forebodies. Ninth International Space Planes and Hypersonic Systems and Technologies Conference and Third Weakly Ionized Gases Workshop, AIAA 99-4816, 1–27, <https://doi.org/10.2514/6.1999-4816>.
- COESA, 1976: *U.S. Standard Atmosphere, 1976*. NOAA, 227 pp.
- Crowther, W. J., and P. J. Lamont, 2000: A neural network approach to the calibration of a flush air data system. University of Manchester School of Engineering Tech. Rep., 21 pp.
- Foresee, F. D., and M. T. Hagan, 1997: Gauss-Newton approximation to Bayesian learning. *The 1997 IEEE International Conference on Neural Networks Proceedings*, Vol. 3, IEEE, 1930–1935, <https://doi.org/10.1109/ICNN.1997.614194>.
- Gracey, W., 1958: Wind-tunnel investigation of a number of total-pressure tubes at high angles of attack: Subsonic, transonic, and supersonic speeds. Annual report of the National Advisory Committee for Aeronautics (43rd): Administrative report including technical report nos. 1296 to 1341, Langley Aeronautical Laboratory National Advisory Committee for Aeronautics Tech. Rep. NACA-TR-1303, 495–504.
- Haering, E. A., 1995: Airdata measurement and calibration. NASA Tech. Memo. 104316, 19 pp.
- Hagan, M. T., and M. B. Menhaj, 1994: Training feedforward networks with the Marquardt algorithm. *IEEE Trans. Neural Networks*, **5**, 989–993, <https://doi.org/10.1109/72.329697>.
- Hornik, K., M. Stinchcombe, and H. White, 1989: Multilayer feedforward networks are universal approximators. *Neural Networks*, **2**, 359–366, [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
- Houston, A. L., B. Argrow, J. Elston, J. Lahowetz, and E. W. Frew, 2012: The Collaborative Colorado–Nebraska Unmanned Aircraft System Experiment. *Bull. Amer. Meteor. Soc.*, **93**, 39–54, <https://doi.org/10.1175/2011BAMS3073.1>.
- , R. Laurence, T. Nichols, S. Waugh, B. Argrow, and C. Ziegler, 2016: Intercomparison of unmanned aircraftborne and mobile mesonet atmospheric sensors. *J. Atmos. Oceanic Technol.*, **33**, 1569–1582, <https://doi.org/10.1175/JTECH-D-15-0178.1>.
- Johansen, E. S., O. K. Rediniotis, and G. Jones, 2001: The compressible calibration of miniature multi-hole probes. *J. Fluids Eng.*, **123**, 128–138, <https://doi.org/10.1115/1.1334377>.
- Kocer, G., M. Mansour, N. Chokani, R. Abhari, and M. Muller, 2011: Full-scale wind turbine near-wake measurements using an instrumented uninhabited aerial vehicle. *J. Solar Energy Eng.*, **133**, 041011, <https://doi.org/10.1115/1.4004707>.
- Larson, T. J., and P. M. Siemers III, 1981: Subsonic tests of an all-flush-pressure-orifice air data system. NASA Tech. Paper NASA TP-1871, 50 pp.
- , S. G. Flechner, and P. M. Siemers III, 1980: Wind tunnel investigation of an all flush orifice air data system for a large subsonic aircraft. NASA Tech. Paper NASA TP-1642, 98 pp.
- , S. A. Whitmore, L. J. Ehernberger, J. B. Johnson, and P. M. Siemers III, 1987: Qualitative evaluation of a flush air data system at transonic speeds and high angles of attack. NASA Tech. Paper NASA TP-2716, 62 pp.
- Laurence, R. J., III, J. S. Elston, and B. M. Argrow, 2015: A low-cost system for wind field estimation through sensor networks and aircraft design. *Proc. AIAA Infotech @ Aerospace*, Kissimmee, FL, AIAA SciTech Forum, AIAA 2015-1425, 13 pp., <https://doi.org/10.2514/6.2015-1425>.
- , B. M. Argrow, and E. W. Frew, 2017: Wind tunnel results for a distributed flush airdata system. *J. Atmos. Oceanic Technol.*, **34**, 1519–1528, <https://doi.org/10.1175/JTECH-D-16-0242.1>.
- MacKay, D. J. C., 1992: Bayesian interpolation. *Neural Comput.*, **4**, 415–447, <https://doi.org/10.1162/neco.1992.4.3.415>.
- Marquardt, D. W., 1963: An algorithm for least-squares estimation of nonlinear parameters. *J. Soc. Ind. Appl. Math.*, **11**, 431–441, <https://doi.org/10.1137/0111030>.
- MathWorks, 2017a: *initnw*: Nguyen–Widrow layer initialization function. Accessed 19 June 2017, <https://www.mathworks.com/help/nnet/ref/initnw.html>.
- , 2017b: *purelin*: Linear transfer function. Accessed 6 November 2017, <https://www.mathworks.com/help/nnet/ref/purelin.html>.
- , 2017c: *tansig*: Hyperbolic tangent sigmoid transfer function. Accessed 6 November 2017, <https://www.mathworks.com/help/nnet/ref/tansig.html>.
- Mohammadi, B., and O. Pironneau, 1994: *Analysis of the K-Epsilon Turbulence Model*. Wiley-Masson Series Research in Applied Mathematics, Book 2, Wiley, 212 pp.
- Moller, M. F., 1993: A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, **6**, 525–533, [https://doi.org/10.1016/S0893-6080\(05\)80056-5](https://doi.org/10.1016/S0893-6080(05)80056-5).
- Nguyen, D., and B. Widrow, 1990: Improving the learning speed of 2-layer neural networks by choosing initial values of the adaptive weights. *1990 IJCNN: International Joint Conference on Neural Networks*, Vol. 3, IEEE, 21–26, <https://doi.org/10.1109/IJCNN.1990.137819>.
- Nichols, T., B. Argrow, and D. Kingston, 2017: Error sensitivity analysis of small UAS wind-sensing systems. *Proc. AIAA*

- Information Systems—AIAA Infotech @ Aerospace, Grapevine, TX, AIAA SciTech Forum, AIAA 2017-0647, 15 pp., <https://doi.org/10.2514/6.2017-0647>.
- Pačes, P., K. Draxler, V. Hanzal, T. Čenský, and O. Vaško, 2010: A combined angle of attack and angle of sideslip smart probe with twin differential sensor modules and doubled output signal. *IEEE Sensors 2010*, IEEE, 284–289, <https://doi.org/10.1109/ICSENS.2010.5689866>.
- Quindlen, J. F., and J. W. Langelaan, 2013: Flush air data sensing for soaring-capable UAVs. *Proc. 51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, Grapevine, TX, AIAA, AIAA 2013-1153, 17 pp., <https://doi.org/10.2514/6.2013-1153>.
- RECUV, 2017: RECUV highlights. University of Colorado Boulder Research and Engineering Center for Unmanned Vehicles, accessed 19 December 2017, <https://www.colorado.edu/recuv/>.
- Reuder, J., M. O. Jonassen, and H. Ólafsson, 2012: The Small Unmanned Meteorological Observer SUMO: Recent developments and applications of a micro-UAS for atmospheric boundary layer research. *Acta Geophys.*, **60**, 1454–1473, <https://doi.org/10.2478/s11600-012-0042-8>.
- Rohloff, T. J., 1998: Air data sensing from surface pressure measurements using a neural network method. *AIAA J.*, **36**, 2094–2101, <https://doi.org/10.2514/2.312>.
- , S. A. Whitmore, and I. Catton, 1999: Fault-tolerant neural network algorithm for flush air data sensing. *J. Aircr.*, **36**, 541–549, <https://doi.org/10.2514/2.2489>.
- Samy, I., I. Postlethwaite, D. Gu, and J. Green, 2010: Neural-network-based flush air data sensing system demonstrated on a mini air vehicle. *J. Aircr.*, **47**, 18–31, <https://doi.org/10.2514/1.44157>.
- Siemens, 2017: STAR-CCM+. Accessed 31 March 2017, <http://mdx.plm.automation.siemens.com/star-ccm-plus>.
- Siemers, P., H. Wolf, and P. F. Flanagan, 1983: Shuttle Entry Air Data System concepts applied to Space Shuttle Orbiter flight pressure data to determine air data—STS 1-4. *Proc. 21st Aerospace Sciences Meeting*, Reno, NV, AIAA, 12 pp., <https://doi.org/10.2514/6.1983-118>.
- TE Connectivity, 2015: MS8607-02BA01, PHT combination sensor. Data Sheet, 22 pp., <http://www.te.com/usa-en/product-CAT-BLPS0018.html>.
- Telionis, D., Y. Yang, and O. Rediniotis, 2009: Recent developments in multi-hole probe (MHP) technology. *Proc. COBEM 2009/20th Int. Congress of Mechanical Engineering*, Gramado, RS, Brazil, Associação Brasileira de Engenharia e Ciências Mecânicas, COB09-3415, 29 pp., <http://abcm.org.br/anais/cobem/2009/pdf/COB09-3415.pdf>.
- Van den Kroonenberg, A., T. Martin, M. Buschmann, J. Bange, and P. Vörsmann, 2008: Measuring the wind vector using the autonomous mini aerial vehicle M²AV. *J. Atmos. Oceanic Technol.*, **25**, 1969–1982, <https://doi.org/10.1175/2008JTECHA1114.1>.
- Vijayagopal, R., M. Pathak, and O. Rediniotis, 1998: Miniature multi-hole pressure probes—Their neural network calibration and frequency response enhancement. *Proc. 36th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, AIAA, 15 pp., <https://doi.org/10.2514/6.1998-204>.
- Whitmore, S. A., 1991: Development of a pneumatic high-angle-of-attack flush airdata sensing (HI-FADS) system. NASA Tech. Memo. 104241, 28 pp.
- , and C. T. Leondes, 1991: Pneumatic distortion compensation for aircraft surface pressure sensing devices. *J. Aircr.*, **28**, 828–836, <https://doi.org/10.2514/3.46105>.
- , T. Moes, and T. Larson, 1990: Preliminary results from a subsonic high-angle-of-attack flush airdata sensing (HI-FADS) system—Design, calibration, algorithm development, and flight test evaluation. *Proc. 28th Aerospace Sciences Meeting*, Reno, NV, AIAA, AIAA-90-0232, 12 pp., <https://doi.org/10.2514/6.1990-232>.
- , —, and C. Leondes, 1992: Failure detection and fault management techniques for flush airdata sensing systems. *Proc. 30th Aerospace Sciences Meeting and Exhibit*, Reno, NV, AIAA, AIAA-92-0263, <https://doi.org/10.2514/6.1992-263>.
- , B. R. Cobleigh, and E. A. Haering, 1997: Design and calibration of the X-33 Flush Airdata Sensing (FADS) System. *Proc. 36th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, NV, AIAA, 36 pp., <https://doi.org/10.2514/6.1998-201>.
- Wildmann, N., M. Hofsäß, F. Weimer, A. Joos, and J. Bange, 2014a: MASC—A small Remotely Piloted Aircraft (RPA) for wind energy research. *Adv. Sci. Res.*, **11**, 55–61, <https://doi.org/10.5194/asr-11-55-2014>.
- , S. Ravi, and J. Bange, 2014b: Towards higher accuracy and better frequency response with standard multi-hole probes in turbulence measurement with remotely piloted aircraft (RPA). *Atmos. Meas. Tech.*, **7**, 1027–1041, <https://doi.org/10.5194/amt-7-1027-2014>.
- Wythoff, B. J., 1993: Backpropagation neural networks: A tutorial. *Chemom. Intell. Lab. Syst.*, **18**, 115–155, [https://doi.org/10.1016/0169-7439\(93\)80052-J](https://doi.org/10.1016/0169-7439(93)80052-J).