https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

Horizontal wind estimation from consumer-drone hover telemetry: comparing machine-learning and calibration approaches

Alexander Polle[1], Lukas Grosch[1], Alexandros Panagiotis Poulidis[1], Mihalis Vrekoussis[1,2,3], and Thorsten Warneke[1]

[1]Institute of Environmental Physics, University of Bremen, 28359 Bremen, Germany
[2]Center for Marine Environmental Sciences (MARUM), University of Bremen, 28359 Bremen, Germany
[3]Climate and Atmosphere Research Center (CARE-C), The Cyprus Institute, 2121 Nicosia, Cyprus

Correspondence: Alexander Polle (polle@uni-bremen.de)

Abstract. Mounting external anemometers on small uncrewed aerial vehicles adds payload weight and complicates flight operations. During hover, horizontal wind can instead be derived from the drone tilt required to hold position. We compare two such tilt-based approaches using flight telemetry from an unmodified DJI Air 3S. Ridge regression fits a linear model to pitch- and roll-derived features and requires reference measurements for training, while a calibration baseline converts mean tilt to wind speed using quadratic curves fitted during separate near-calm flights and requires no reference anemometer during deployment. Both methods were evaluated across 50 hover sessions in Bremen, Germany, against a reference anemometer 1.6 m above ground. The sessions yielded 71,749 overlapping 30 s windows. Session-mean wind speeds ranged from 0.8 to 7.6 $\mathrm{m\,s^{-1}}$, and the highest 30 s vector-mean speed was 11.5 $\mathrm{m\,s^{-1}}$. In leave-one-flight-out cross-validation, Ridge had a speed root-mean-square error (RMSE) of 0.32 $\mathrm{m\,s^{-1}}$ and a direction mean absolute error (MAE) of 7.0°, compared with 0.39 $\mathrm{m\,s^{-1}}$ and 8.4° for the baseline. For session-mean winds above 2 $\mathrm{m\,s^{-1}}$, direction MAE was 6.0° for Ridge and 7.0° for the baseline, close to the estimated 5–7° directional uncertainty of the reference setup. Across the six sessions above 6 $\mathrm{m\,s^{-1}}$, Ridge and the baseline had similar speed RMSEs of 0.49 and 0.48 $\mathrm{m\,s^{-1}}$. Tree-based models (XGBoost and LightGBM) did not improve upon Ridge. Hover telemetry can therefore provide accurate horizontal wind estimates without dedicated payload sensors. Further validation is needed for other airframes, greater heights, and sustained stronger winds.

1 Introduction

Small uncrewed aerial vehicles (UAVs) are used to measure winds in the atmospheric boundary layer, but many systems rely on external anemometers such as sonic or multi-hole probes (Reuter et al., 2021; Meier et al., 2022). These instruments add weight and cost, require power, and make field operations more complex. Added payload can affect regulatory classification (European Union, 2019) and may alter the aerodynamic response, stability, or endurance of the aircraft (Ghirardelli et al., 2025; Soltaninezhad et al., 2025). External sensors are also exposed to rotor-induced airflow (Yang et al., 2025) and may require dedicated calibration flights (Reuter et al., 2021). Boom-mounted sonic anemometers are therefore usually positioned clear of the rotor downwash (Thielicke et al., 2021). Wind estimated from telemetry the aircraft already records requires no additional payload.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

To maintain a fixed position in a moving air mass, a multicopter tilts to counteract wind-induced drag forces. By measuring the pitch and roll response, the horizontal wind vector can be estimated (Palomaki et al., 2017). Early implementations relied on analytical models (Neumann and Bartholmai, 2015; Palomaki et al., 2017). Later work derived wind from the full flight dynamics of the aircraft (González-Rocha et al., 2019) or used flight data to establish drag models and platform calibrations (Hattenberger et al., 2022; Bramati et al., 2024). Machine-learning approaches have also been used to represent nonlinear relationships between wind and airframe response (Zhu et al., 2025; García-Gascón et al., 2026). These include nearest-neighbour methods and recurrent networks (Crowe et al., 2020), Gaussian-process regression (Zimmerman et al., 2022b), and random forests (Chen et al., 2026). Using two consumer drones in hover, Crowe et al. (2020) obtained a speed RMSE of 0.61 $\mathrm{m\,s^{-1}}$. Meier et al. (2022) applied dynamic models to DJI Phantom 4 flights in Switzerland and Norway. Reuter et al. (2021) found greater scatter in attitude-based estimates than in measurements from a calibrated anemometer carried on the same flight. To our knowledge, direct comparisons between a reference-trained regression model and an anemometer-free translational calibration on the same consumer multicopter across multiple wind regimes have not been reported.

This study assesses the accuracy with which an unmodified DJI Air 3S can estimate horizontal wind from its internal hover telemetry. Two attitude-based methods with different data requirements are compared. Ridge regression applies a linear model to features derived from pitch and roll, with an L_2 penalty that stabilizes the coefficients when these features are correlated (Hoerl and Kennard, 1970). The model must be trained against measurements from a reference anemometer. The calibration baseline instead estimates the wind components from mean pitch and roll using quadratic curves fitted to near-calm flight data. Once calibrated, it can be used without a reference anemometer.

Both methods were evaluated across 50 hover sessions using leave-one-flight-out validation, which held out complete sessions rather than neighbouring windows. In each fold, Ridge was trained on the remaining sessions using the reference wind measurements. The baseline coefficients came from separate near-calm flights in which aircraft-reported ground speed was paired with tilt along each body axis. Two gradient-boosted tree models were also evaluated using the same features and validation folds as Ridge.

2 Methods

2.1 Platform and data collection

Figure 1 summarizes the data streams and validation workflow. Measurements were collected during 50 hover sessions at an open field and a hilltop in Bremen, Germany, between 1 October and 18 November 2025 (Supplement Sect. S1.3). Flights used a DJI Air 3S weighing 724 g and running firmware v01.00.1500 (DJI, 2025). The aircraft maintained position using its onboard global navigation satellite system (GNSS) and inertial measurement unit (IMU), and recorded attitude telemetry throughout. An LI-550 TriSonica Mini ultrasonic anemometer measured the reference wind (LI-COR, 2025). A custom data logger recorded it (Polle, 2026). The sensor head was set 1.6 m above ground, and the aircraft flew at the same nominal height with at least 2 m horizontal separation (Fig. 2a), a distance at which no downwash effect was detectable (Supplement Sect. S1.15). The reference setup has a conservative directional uncertainty range of 5–7° (Sect. 4.2).

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

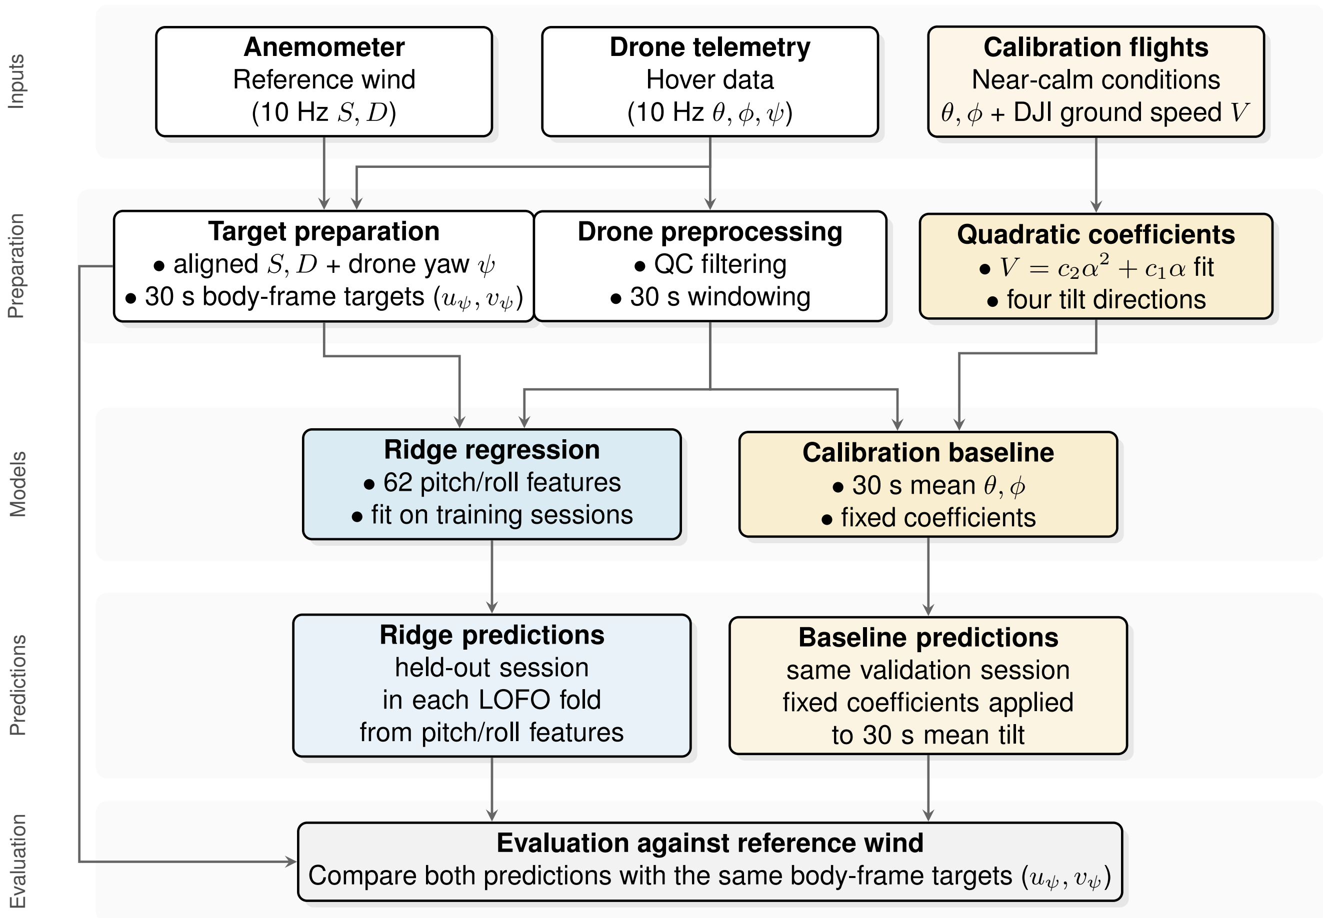

Figure 1. Workflow for the Ridge and calibration-baseline analyses. White boxes show the steps shared by both methods, while blue boxes show the Ridge branch and orange boxes the baseline branch. In each leave-one-flight-out fold, Ridge is fitted on 49 sessions using their reference targets and then predicts the remaining session. The calibration baseline instead applies fixed coefficients from separate near-calm flights to that same session. Both sets of predictions are evaluated against the same 30 s body-frame reference targets (Sect. 2.6). XGBoost and LightGBM use the Ridge features and validation folds but are omitted for clarity.

Each session lasted approximately 30 min. At takeoff, the aircraft was placed downwind of the reference sensor and then hovered in place for the rest of the session. Each session was divided into 4–6 min segments flown at different headings, varying the headwind and crosswind components relative to the aircraft. This pattern covered all four drone-relative quadrants (Fig. 2b), with the share of windows ranging from 21.9 % in the headwind sector to 30.0 % in the left crosswind sector. After filtering and windowing, the 50 sessions yielded 71,749 overlapping 30 s windows (Sect. 2.3). Session-mean wind speeds ranged from 0.8 to 7.6 m s^{-1}, and the highest 30 s vector-mean speed was 11.5 m s^{-1}.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

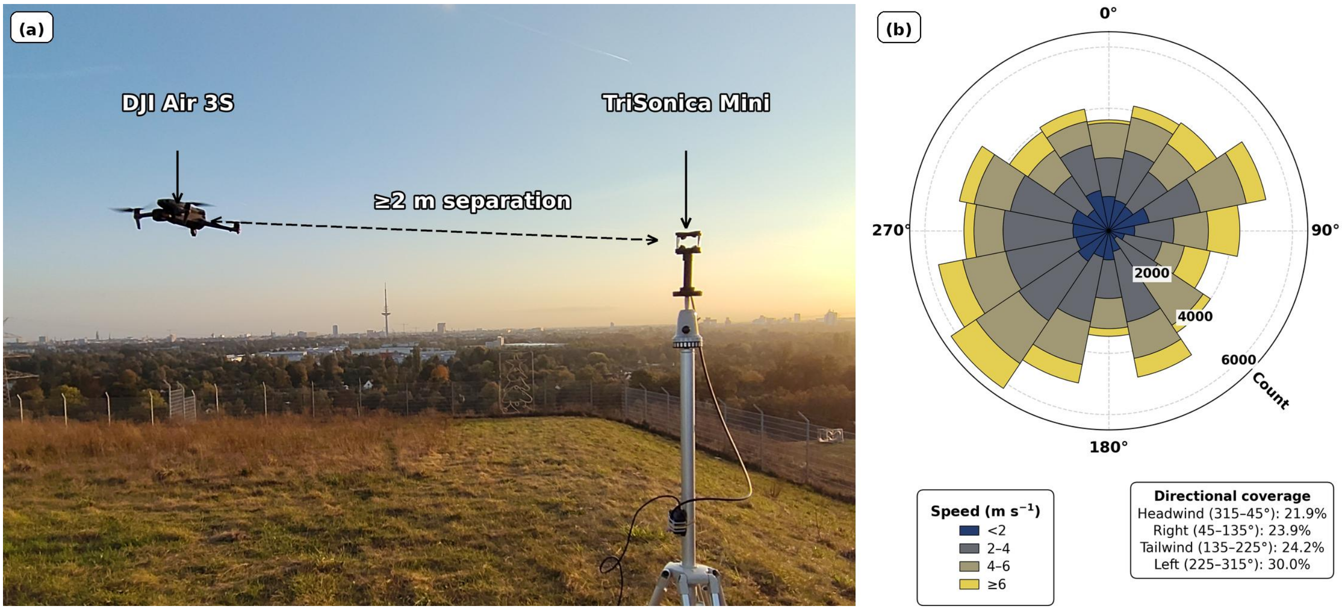

Figure 2. Experimental setup and dataset coverage. (a) DJI Air 3S and TriSonica Mini reference sensor at the hilltop site in Bremen, Germany. The sensor head was set 1.6 m above ground. The aircraft flew at the same nominal height with at least 2 m horizontal separation. (b) Distribution of drone-relative wind direction across all 71,749 rolling 30 s windows with a 1 s step. A direction of 0° denotes headwind, 90° right crosswind, 180° tailwind, and 270° left crosswind. Colours indicate 30 s vector-mean wind-speed intervals in m s^{-1}, and the inset summarizes directional coverage.

2.2 Data extraction, time alignment, and reference targets

The controller stored encrypted flight telemetry, which was decrypted after each flight using `dji-log-parser` (Vauvillier, 2025). Pitch and roll provided the tilt measurements, yaw provided the aircraft heading, and deflections of the `RC.rudder` stick identified intentional yaw manoeuvres during quality control.

The drone and reference timestamps were converted to Coordinated Universal Time (UTC), and both series were interpolated linearly onto a common 10 Hz grid over their overlapping period. Before interpolation, yaw and wind direction were unwrapped to avoid discontinuities at the 0°–360° boundary, while pitch and roll were interpolated directly. No additional timing offset was estimated or applied.

A separate post-analysis timing diagnostic cross-correlated the held-out Ridge and reference wind-speed traces. The peak relative lag ranged from -1 to $+4$ s in 49 sessions. One session had a lag magnitude of 38 s. No time shift was applied on the basis of this diagnostic (Supplement Sect. S1.1).

Because pitch and roll are measured along the longitudinal and lateral axes of the drone, the same geographic wind projects differently onto these axes as the aircraft heading changes. The aligned horizontal wind vector, represented by speed S and meteorological “from” direction D, was therefore rotated by the aircraft yaw ψ into the drone body frame (Zimmerman et al.,

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

2022a):

$$D_\psi = D - \psi,$$
$$u_\psi = S\cos D_\psi, \qquad (1)$$
$$v_\psi = S\sin D_\psi.$$

Here D_ψ is the drone-relative wind direction. A value of $0°$ denotes a headwind and $180°$ a tailwind. The longitudinal component u_ψ is positive for a headwind, and the lateral component v_ψ is positive for wind from the right. Directions are stored and reported in degrees, but they are converted to radians inside the sine and cosine functions.

The body-frame targets remain aligned with the pitch and roll axes as the aircraft turns, allowing a single relationship to be fitted without treating each absolute heading separately. A mobile compass operating in true-north mode was used to align the anemometer, while aircraft heading came from the onboard yaw estimate (Supplement Sect. S1.2). Any remaining misalignment is included in the directional uncertainty budget (Sect. 4.2). Predicted wind speed $\widehat{S}$ and drone-relative direction $\widehat{D}_\psi$ are recovered by inverting Eq. (1). Because the same yaw rotation applies to the predicted and reference vectors, rotating both back to geographic coordinates would not change their angular difference.

2.3 Data filtering and windowing

Each aligned 10 Hz time series was divided into rolling 30 s windows with a 1 s step, so adjacent windows overlap by 29 s and are not statistically independent. This averages over rapid tilt adjustments while preserving the lower-frequency wind signal. The mean headwind (u_ψ) and crosswind (v_ψ) in each window are the prediction targets.

The calibration baseline uses only the mean pitch and roll of each window, whereas Ridge and the tree models use 62 statistical and dynamical features derived from the same samples (Sect. 2.5). A sensitivity analysis compared 10, 30, and 60 s windows at common endpoints. Errors decreased with longer windows, but the 30 s interval was retained to limit the loss of temporal resolution (Supplement Sect. S1.6).

Quality control (QC) masked the first and last 10 s of each session and periods of intentional yaw manoeuvres identified from `RC.rudder` stick deflection. Continuous intervals shorter than 30 s after these exclusions were also masked. A window was discarded if it overlapped a masked sample, contained a negative reference-speed value used to mark invalid data, or had finite values for fewer than 80 % of its samples (Table S1). No additional position or altitude filter was applied. The first five retained windows of each session were removed because they lacked the history required by the 5 s lag features. This left 71,749 windows for analysis.

2.4 Calibration baseline

The calibration baseline maps tilt to airspeed using four quadratic curves for forward pitch, backward pitch, left roll, and right roll. The coefficients in Table 1 were derived from calibration flights at an open field in near-calm conditions, where the mean wind was below $0.5\ \mathrm{m\,s^{-1}}$. The aircraft flew at speed steps from 1 to 13 $\mathrm{m\,s^{-1}}$ along each body axis, with three repetitions for every speed and direction. Steady samples from the three repetitions were pooled into one calibration point, giving 13 points

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

for each curve. The fits used the ground speed reported by the drone during steady horizontal flight. Pairing legs flown in opposite geographic directions reduced the influence of the remaining ambient wind (Fig. 3). Any scale error in the reported ground speed would remain in the baseline estimates.

The calibration relation is

$$V = c_2\alpha^2 + c_1\alpha, \tag{2}$$

where V is the speed magnitude in $\mathrm{m\,s^{-1}}$ and the signed angle α is either pitch θ or roll ϕ in degrees. During fitting, V was the reported ground speed, which approximated airspeed under near-calm conditions. When applied during hover, the fitted relation estimates airspeed from tilt.

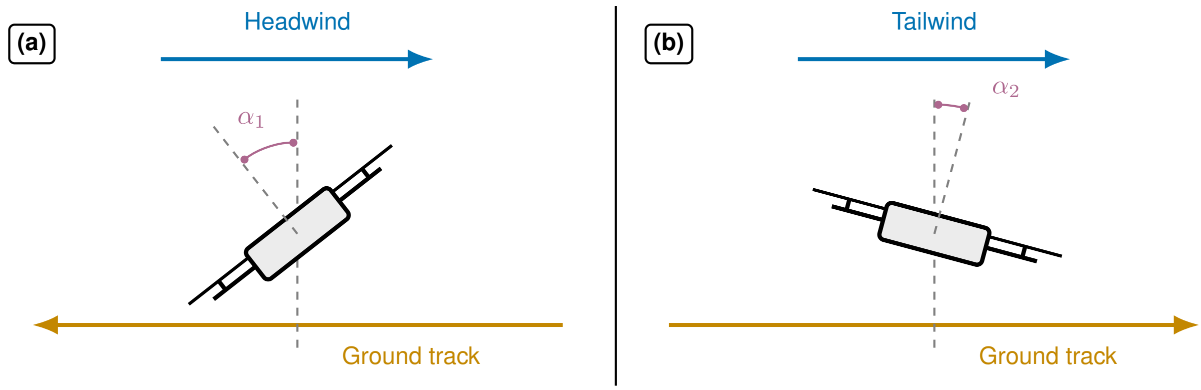

Figure 3. Schematic showing how opposing calibration legs reduce the effect of the remaining ambient wind. (a) The drone flies against the wind and requires tilt α_1 to maintain the selected ground speed. (b) After a 180° heading change, it flies with the wind and requires tilt α_2. At each speed step, steady samples from matched legs in opposite geographic directions were pooled before fitting.

The quadratic relation was fitted directly to the calibration points. Previous studies calibrated tilt in a wind tunnel (Neumann and Bartholmai, 2015) or during hover against a reference anemometer (Palomaki et al., 2017), whereas the baseline used here relies on aircraft-reported ground speed during near-calm translation. The baseline can be applied without a reference instrument once calibrated. Using the curves in stationary hover assumes that the relation measured in translational flight also holds there. Figure 4 shows the four fitted curves.

During steady hover, the baseline model converts the 30 s average pitch θ and roll ϕ within each window into headwind $\hat{u}_\psi$ and crosswind $\hat{v}_\psi$ components according to

$$\hat{u}_\psi = \begin{cases} |f_{\mathrm{fwd}}(\theta)| & \theta < 0 \\ -|f_{\mathrm{bwd}}(\theta)| & \theta \geq 0 \end{cases}, \quad \hat{v}_\psi = \begin{cases} -|f_{\mathrm{left}}(\phi)| & \phi < 0 \\ |f_{\mathrm{right}}(\phi)| & \phi \geq 0 \end{cases} \tag{3}$$

where each f_{axis} is the corresponding sign-dependent quadratic function from Eq. (2) with the coefficients in Table 1. A forward tilt ($\theta < 0$) maps to a headwind ($\hat{u}_\psi > 0$), and a backward tilt ($\theta > 0$) maps to a tailwind ($\hat{u}_\psi < 0$). A left roll ($\phi < 0$)

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

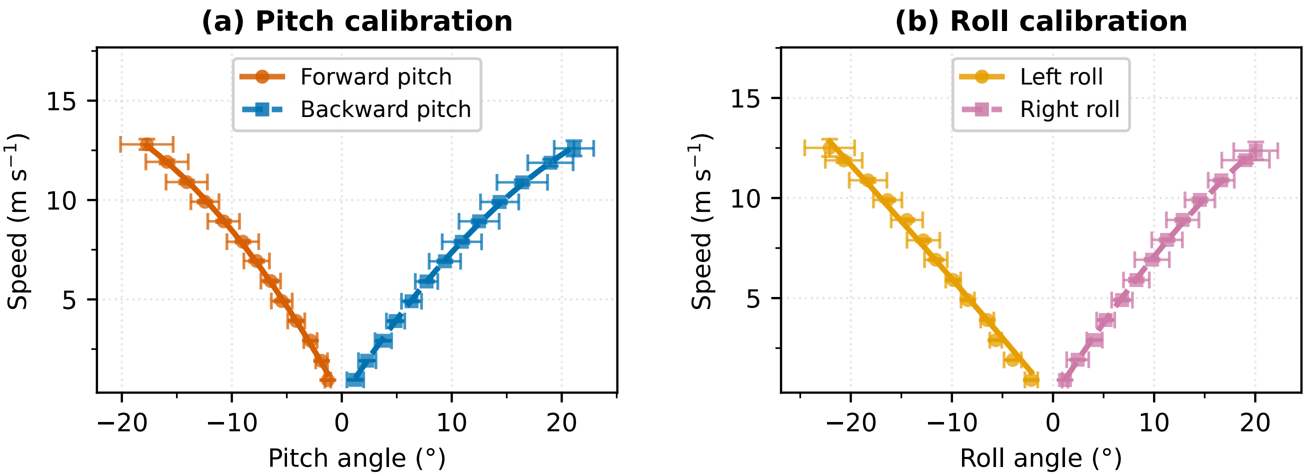

Figure 4. Tilt–airspeed calibration results. (a) Longitudinal ground-speed magnitude versus pitch angle. (b) Lateral ground-speed magnitude versus roll angle. Each point pools the steady samples from three repetitions at one speed and body-axis direction. Error bars show ± 1 standard deviation across those samples. Curves show the zero-intercept quadratic fits of Eq. (2) for negative and positive tilt angles across the 1–13 $\mathrm{m\,s^{-1}}$ calibration range.

Table 1. Quadratic coefficients for the tested DJI Air 3S unit in the calibration relation $V = c_2\alpha^2 + c_1\alpha$. Forward pitch and left roll correspond to negative angles, while backward pitch and right roll correspond to positive angles. Each curve was fitted by unweighted least squares to 13 calibration points with the intercept fixed at zero, leaving 11 degrees of freedom. The $\pm$ values give one standard error of each coefficient. RMSE is calculated from the 13 fit residuals.

Direction	c_2 ($\mathrm{m\,s^{-1}\,deg^{-2}}$)	c_1 ($\mathrm{m\,s^{-1}\,deg^{-1}}$)	**RMSE** ($\mathrm{m\,s^{-1}}$)
Pitch forward	-0.0168 ± 0.0010	-1.016 ± 0.014	0.128
Pitch backward	-0.01225 ± 0.00041	$+0.8595 \pm 0.0069$	0.075
Roll left	-0.0013 ± 0.0014	-0.609 ± 0.025	0.297
Roll right	-0.00871 ± 0.00063	$+0.795 \pm 0.010$	0.106

maps to wind from the left ($\hat{v}_\psi < 0$), and a right roll ($\phi > 0$) maps to wind from the right ($\hat{v}_\psi > 0$). At zero pitch or roll, the corresponding estimated wind component is zero. Observed window-mean pitch angles ($-11.0°$ to $+14.9°$) and roll angles ($-17.5°$ to $+19.3°$) remained within the maximum calibrated tilt magnitudes (17.7° and 21.1° in pitch, 22.1° and 20.1° in roll), so the baseline never exceeded its upper calibration limits. Near zero tilt, where calibration flights began at ~ 1 $\mathrm{m\,s^{-1}}$ (1.2°–2.1° tilt), the baseline relies on the physically constrained zero intercept.

The left-roll curve is almost linear and has the largest residual RMSE of the four fits. Its quadratic coefficient is smaller in magnitude than its standard error, so the calibration does not resolve curvature in this direction (Table 1). A linear fit of

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

roll angle against signed ground speed across both roll directions has an intercept of $-0.8°$. This offset may contribute to the asymmetry, but the data do not establish its cause. The zero-intercept baseline does not correct this offset and also treats pitch and roll as independent axes. Supplement Sect. S1.20 examines whether this assumption increases error under diagonal winds.

2.5 Ridge regression model and features

Each 30 s hover window was represented by 62 features derived from pitch and roll. These features describe mean tilt, variation within the window, and changes over recent windows, together with nonlinear terms and products between the two axes. Ridge regression (Hoerl and Kennard, 1970) maps these features to the headwind and crosswind components in the drone body frame. Table S2 gives their complete definitions.

The feature definitions were based on physical reasoning and an initial examination of the first 34 campaign sessions, collected from 1 to 10 October 2025. Ridge was implemented in scikit-learn version 1.7.2 with a regularization strength of $\lambda = 1.0$, corresponding to `alpha`=1.0 (Pedregosa et al., 2011). The feature set and regularization strength were then kept fixed in every validation fold. Because the first 34 sessions remain in the final dataset, cross-validation is not an external test of the feature design.

The 10 Hz pitch and roll samples in each window were summarized by their means, standard deviations, extrema, absolute values, and squared values, together with the mean product of the two axes. Temporal features describe changes between retained windows and trends over preceding windows. Supplement Sect. S1.5 defines these features and tests the effect of resetting them at QC gaps. Yaw was excluded because it describes aircraft heading rather than the aerodynamic response to horizontal wind.

For each fold, the features were standardized to zero mean and unit variance using statistics from the training sessions only. Ridge is linear in the 62 engineered features rather than in the raw pitch and roll values. Squared and interaction terms therefore allow the fitted response to include curvature and cross-axis relationships. The L_2 penalty helps stabilize the coefficients when predictors are correlated. An unregularized least-squares sensitivity test changed the reported metrics by less than 0.001 $\mathrm{m\,s^{-1}}$ and $0.1°$ (Supplement Sect. S1.8). Separate models were fitted for u_ψ and v_ψ.

For comparison, the gradient-boosted tree ensembles XGBoost and LightGBM were trained with the same features and folds, using settings fixed beforehand (Supplement Sect. S1.10 and Sect. 3.5). This comparison tests whether the tree models improve accuracy when given the same inputs. Ridge was selected as the primary model after the comparison, so its cross-validated errors do not include the uncertainty introduced by model selection.

2.6 Validation protocol and evaluation metrics

Validation used leave-one-flight-out (LOFO) cross-validation across the 50 hover sessions, holding out all windows from one session while Ridge was trained on the other 49. Because neighbouring 30 s windows overlap by 29 s, splitting individual windows would place closely related observations in both the training and validation sets (Roberts et al., 2017). Metrics were therefore calculated for each held-out session and then averaged, giving every session equal weight. The calibration baseline had fixed coefficients and was evaluated on the same windows.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

Two additional analyses used broader holdouts. Leave-three-flights-out (L3FO) validation evaluated Ridge on all 19,600 combinations of three held-out sessions. Because individual sessions recur across combinations, the resulting spread is not an independent uncertainty estimate (Supplement Sect. S1.11). A second analysis held out all sessions from one measurement day at a time, preventing data from the same day from appearing in both training and validation (Supplement Sect. S1.22).

Wind speed root-mean-square error (RMSE) was calculated from the magnitude of the 30 s mean horizontal wind vector rather than from the mean of the scalar 10 Hz speeds. Direction mean absolute error (MAE) was calculated from the absolute circular difference between predicted and reference directions. Because wind direction is ill-conditioned at low speeds, direction MAE is also reported for a stronger-wind subset. For LOFO and the calibration baseline, this subset contains sessions with a mean reference wind above 2 $\mathrm{m\,s^{-1}}$. For L3FO, the same threshold is applied to the mean of each held-out three-session combination. Section 3.3 reports window-level results, and Supplement Sect. S1.7 gives the metric equations.

3 Results

3.1 Aggregate performance

In LOFO validation, Ridge has a speed RMSE of 0.32 $\mathrm{m\,s^{-1}}$ and a direction MAE of 7.0°, compared with 0.39 $\mathrm{m\,s^{-1}}$ and 8.4° for the calibration baseline (Table 2). When the direction metric is restricted to sessions with mean winds above 2 $\mathrm{m\,s^{-1}}$, the MAE falls to 6.0° for Ridge and 7.0° for the baseline. The mean signed direction errors in this subset are small (Sect. 4.2).

The body-frame vector RMSE is 0.55 $\mathrm{m\,s^{-1}}$ for Ridge and 0.67 $\mathrm{m\,s^{-1}}$ for the baseline, while their session-averaged speed biases are -0.02 and $+0.18$ $\mathrm{m\,s^{-1}}$ (Table 2). Among individual sessions, Ridge gives the lower speed RMSE in 82 % of cases and the lower direction MAE in 72 %. Section 4.1 reports the paired comparison, and Table S4 gives the L3FO sensitivity results.

In one high-wind session, the telemetry-based estimates and reference trace exhibit a relative timing offset (Supplement Sect. S1.14 and Fig. S1). The session was retained in the primary evaluation, so the reported metrics include this offset. Excluding it lowers the aggregate speed RMSE to 0.30 $\mathrm{m\,s^{-1}}$ for Ridge and 0.38 $\mathrm{m\,s^{-1}}$ for the calibration baseline.

3.2 Feature contributions

Because the features were standardized to zero mean and unit variance, coefficient magnitudes can be compared within each fitted component model (Fig. 5). A coefficient gives the change in the predicted wind component, in $\mathrm{m\,s^{-1}}$, for a one-standard-deviation change in its feature. In both models, the three largest absolute coefficients belong to squared tilt magnitude, mean squared roll, and mean squared pitch. These nonlinear terms allow the response to vary with tilt magnitude.

Other large coefficients mainly come from the tilt axis associated with each wind component. The headwind (u_ψ) model uses mean pitch, its 2 and 5 s lags, and the pitch minima and maxima. The crosswind (v_ψ) model uses the corresponding roll terms, with signs following the convention in Sect. 2.2. These patterns describe the fitted model rather than causal effects. Because the L_2 penalty distributes weight across correlated features, no coefficient can be interpreted as an independent contribution.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

Table 2. Session-level performance on 30 s windows, reported as mean ± 1 standard deviation across sessions. Each session has equal weight. Ridge values are held-out LOFO predictions, while the calibration baseline uses fixed coefficients on the same windows. Speed and direction biases are signed prediction-minus-reference errors. Direction bias is wrapped circularly, and vector RMSE combines the two body-frame component errors. The > 2 m s^{-1} rows include sessions whose mean wind exceeds that threshold (Sect. 2.6). Bold values mark the lower error in each row. Window-level results are given in Sect. 3.3.

Metric	Ridge (LOFO)	Calibration baseline
Speed RMSE (m s^{-1})	$\mathbf{0.32 \pm 0.13}$	0.39 ± 0.15
Speed bias (m s^{-1})	$\mathbf{-0.02 \pm 0.15}$	$+0.18 \pm 0.17$
Vector RMSE (m s^{-1})	$\mathbf{0.55 \pm 0.28}$	0.67 ± 0.27
Direction MAE (°)	$\mathbf{7.0 \pm 4.9}$	8.4 ± 5.0
Direction MAE (°), for >2 m s^{-1}	$\mathbf{6.0 \pm 3.7}$	7.0 ± 3.8
Direction bias (°), for >2 m s^{-1}	$\mathbf{+0.4 \pm 6.5}$	-0.7 ± 6.6

Table S2 lists all 62 features, and Table S6 summarizes the selected feature families. A Ridge model restricted to six primary tilt features produced errors close to those of the full model (Supplement Sect. S1.9).

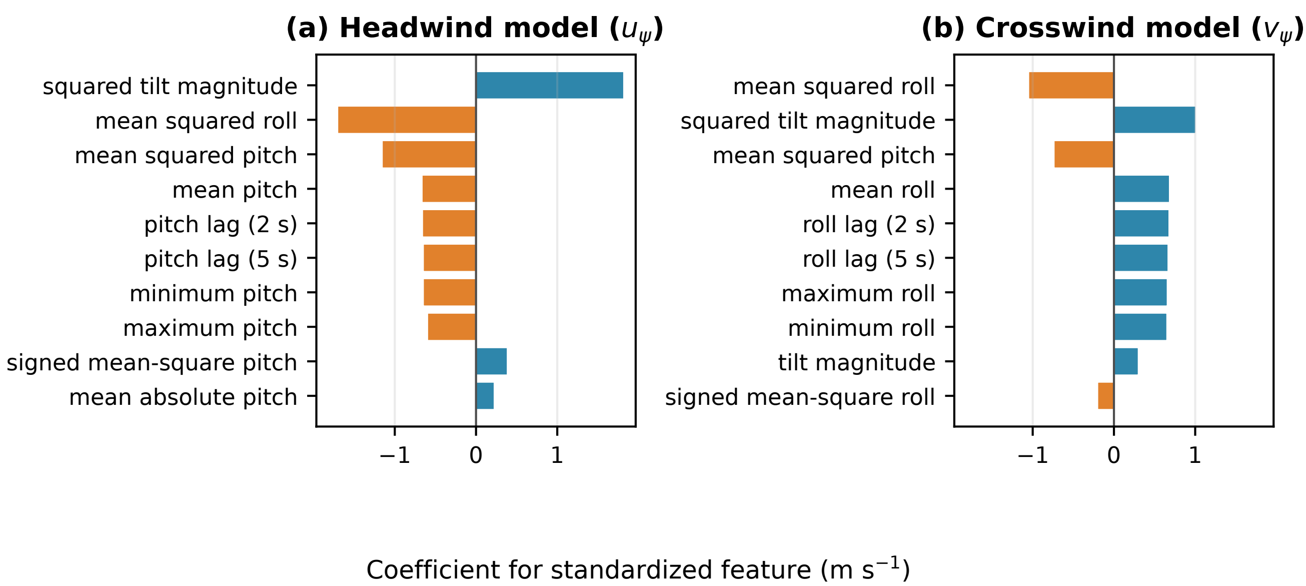

Figure 5. Ridge coefficients fitted across all 50 hover sessions after the features were standardized to zero mean and unit variance. Panels show (a) the headwind (u_ψ) model and (b) the crosswind (v_ψ) model. Each panel contains the ten largest absolute coefficients, with positive values in blue and negative values in orange. The targets retain their original units, so each coefficient is expressed in m s^{-1} per feature standard deviation. Correlated terms share weight through the L_2 penalty and are not independent effects. Table S2 defines the features.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

3.3 Window-level performance

Predictions from both methods cluster around the 1:1 line for wind speed and drone-relative direction (Fig. 6). Across all 71,749 pooled windows, speed RMSE is 0.35 $\mathrm{m\,s^{-1}}$ for Ridge and 0.43 $\mathrm{m\,s^{-1}}$ for the calibration baseline, slightly higher than the session-weighted values in Table 2. For windows with reference speeds above 2 $\mathrm{m\,s^{-1}}$, direction MAE is 5.7° for Ridge and 6.7° for the baseline. Above 5 $\mathrm{m\,s^{-1}}$, Ridge underestimates reference speed by an average of 0.13 $\mathrm{m\,s^{-1}}$, while the baseline overestimates it by 0.14 $\mathrm{m\,s^{-1}}$. These pooled values mask differences in bias among sessions.

3.4 Performance across wind regimes

Performance differs among the three wind-speed regimes (Tables 3 and S7, Fig. S2). Of the 50 sessions, 44 had mean winds below 6 $\mathrm{m\,s^{-1}}$ and six exceeded that threshold. Ridge has lower speed RMSE and direction MAE in both the low- and medium-wind regimes, with its lowest errors occurring between 2 and 6 $\mathrm{m\,s^{-1}}$. Across the 34 sessions in this range, Ridge reduces speed RMSE by a median of 0.09 $\mathrm{m\,s^{-1}}$ relative to the calibration baseline (Fig. 7).

Above 6 $\mathrm{m\,s^{-1}}$, this advantage disappears, with similar RMSEs of 0.49 $\mathrm{m\,s^{-1}}$ for Ridge and 0.48 $\mathrm{m\,s^{-1}}$ for the calibration baseline. Removing the relative-lag session reduces these values to 0.37 and 0.36 $\mathrm{m\,s^{-1}}$ across the remaining five sessions. In the same subset, XGBoost has an RMSE of 0.55 $\mathrm{m\,s^{-1}}$ and LightGBM 0.50 $\mathrm{m\,s^{-1}}$ (Tables S8 and S9). Figure 8 shows a different high-wind session in which the retrievals follow the same temporal changes without a visible phase offset.

Table 3. Model performance across wind speed regimes. Speed RMSE in $\mathrm{m\,s^{-1}}$ and direction MAE in degrees are averaged across the sessions in each regime. n denotes the number of sessions. Lower displayed errors are shown in bold, while equal rounded values are left unmarked.

Wind regime	Model	n	Speed RMSE	Direction MAE
Low ($< 2\ \mathrm{m\,s^{-1}}$)	Ridge	10	**0.32**	**11.0**
	Cal. baseline	10	0.40	13.8
Medium (2–6 $\mathrm{m\,s^{-1}}$)	Ridge	34	**0.29**	**5.9**
	Cal. baseline	34	0.38	7.1
High ($> 6\ \mathrm{m\,s^{-1}}$)	Ridge	6	0.49	6.5
	Cal. baseline	6	**0.48**	6.5

3.5 Comparison with alternative model families

XGBoost version 3.0.5 (Chen and Guestrin, 2016) and LightGBM version 4.6.0 (Ke et al., 2017) were evaluated with the same features and folds as Ridge. Their settings were fixed and were not tuned on the validation data (Supplement Sect. S1.10). Both tree models can represent nonlinear relationships in tabular data (Grinsztajn et al., 2022). With the same inputs, all three trained models have lower aggregate errors than the calibration baseline (Table 4). Ridge has the lowest session-averaged errors, with

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

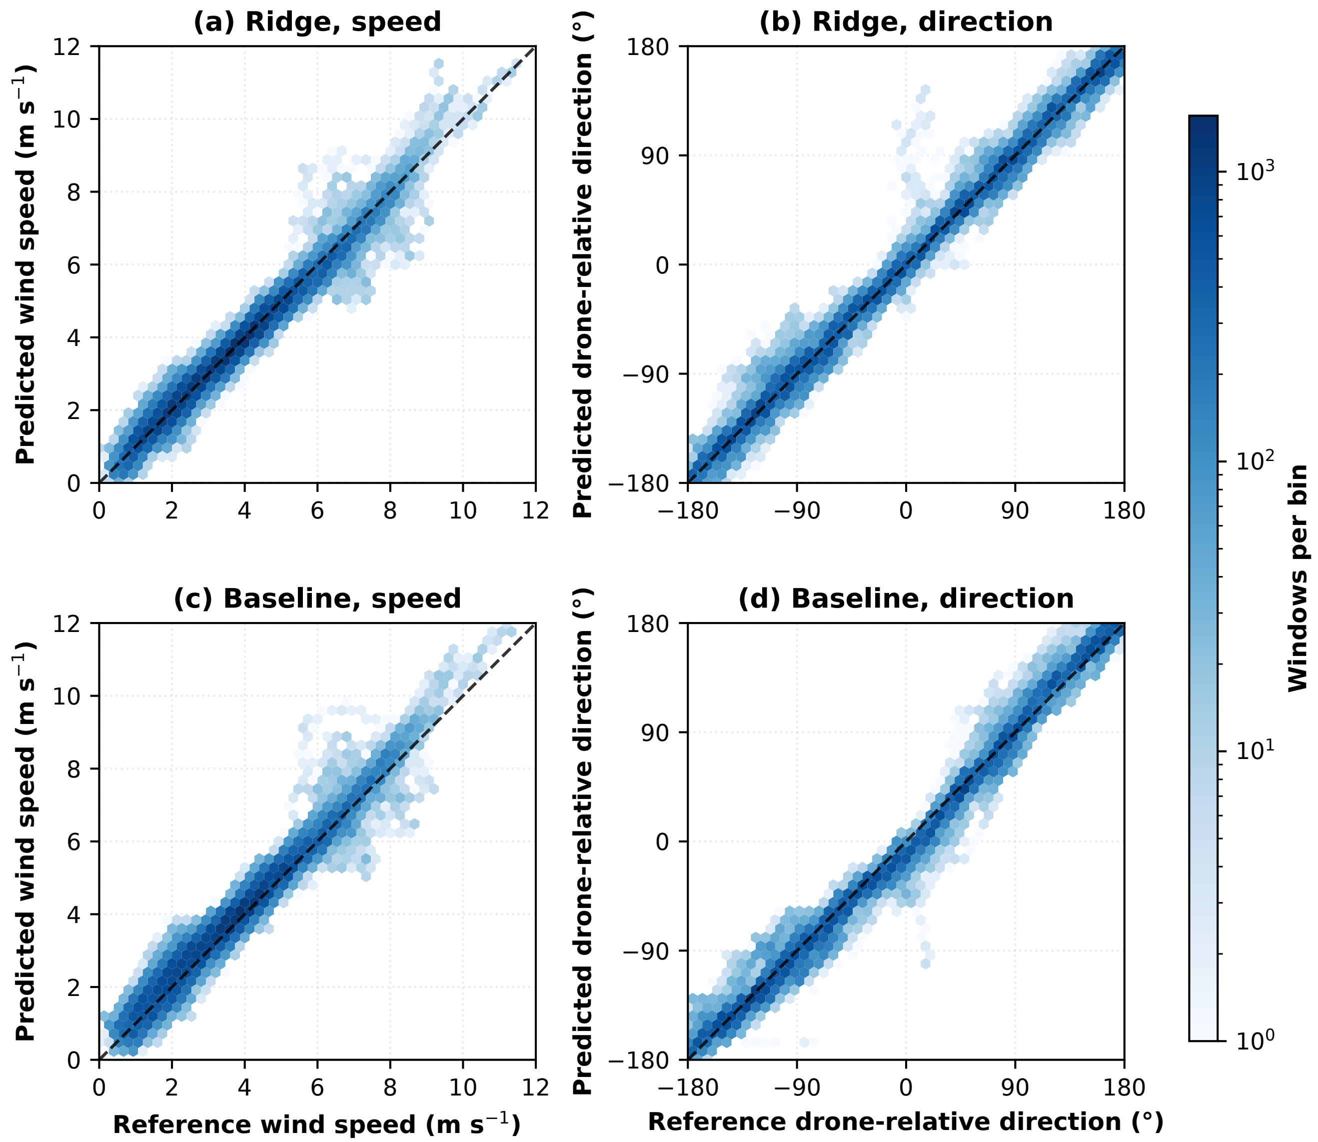

Figure 6. Window-level density of predicted versus reference values across all 71,749 30 s windows, shown on a logarithmic colour scale. Adjacent windows overlap by 29 s and are not statistically independent. (a–b) Ridge model held-out LOFO predictions for (a) wind speed and (b) drone-relative direction. (c–d) Calibration baseline predictions for (c) wind speed and (d) drone-relative direction using fixed coefficients (Table 1). Dashed lines mark the 1:1 reference. The off-diagonal cluster in panels (b) and (d) stems from the single near-calm session with a median wind of 0.7 m s^{-1}, where direction is ill-conditioned. The elevated scatter near 5–9 m s^{-1} in panels (a) and (c) originates primarily from the retained relative-lag session (Sect. 3.1).

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

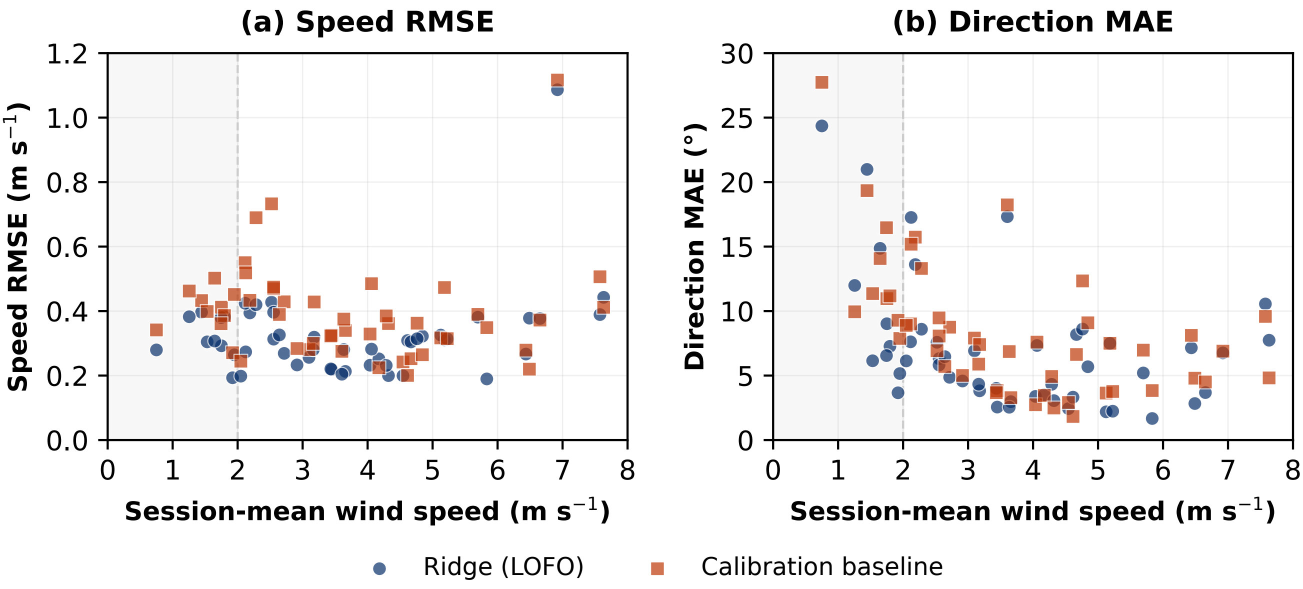

Figure 7. Per-session errors versus session-mean wind speed across the 50 validation hover sessions: (a) speed RMSE and (b) direction MAE for the Ridge model (LOFO, circles) and calibration baseline (squares). Grey shading marks the weak-wind regime below 2 m s^{-1} where wind direction is ill-conditioned. The two elevated points near 6.9 m s^{-1} in panel (a) correspond to the relative-lag session (Sect. 3.1).

Table 4. Performance of the machine-learning models and calibration baseline with equal weight given to each session. The machine-learning models used the same 62-feature matrix and leave-one-flight-out folds. The fixed calibration baseline was evaluated on the same windows. The final column reports speed RMSE for the six sessions with mean wind above 6 m s^{-1}. Bold values mark the lowest error per column.

Model	Speed RMSE (m s^{-1})	Direction MAE (°)	High-wind speed RMSE (m s^{-1})
Ridge regression	**0.32**	**7.0**	0.49
LightGBM (gradient-boosted trees)	0.33	7.4	0.58
XGBoost (gradient-boosted trees)	0.33	7.5	0.63
Calibration baseline	0.39	8.4	**0.48**

a speed RMSE of 0.32 m s^{-1} and a direction MAE of 7.0°. XGBoost and LightGBM both have a speed RMSE of 0.33 m s^{-1}, with direction MAEs of 7.5° for XGBoost and 7.4° for LightGBM.

In the six sessions with mean winds above 6 m s^{-1}, the calibration baseline and Ridge give lower speed errors than the tree models (Tables S8 and S9). That subset is small, so the ranking is not firmly established. Direction MAE is similar across all four methods and ranges from 6.4 to 6.7° in these sessions.

The strongest winds are sparsely represented in the dataset, with only 2.1 % of the 71,749 windows exceeding 8 m s^{-1} and 0.14 % exceeding 10 m s^{-1}. Gradient-boosted trees form piecewise-constant functions and generally do not extrapolate

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

well beyond the feature combinations represented in their training data. In this study, their response at high wind speeds is compressed relative to Ridge and the calibration baseline. The 334 windows above 9 $\mathrm{m\,s^{-1}}$ have a mean reference speed of 9.8 $\mathrm{m\,s^{-1}}$, compared with predicted means of 8.3 $\mathrm{m\,s^{-1}}$ for XGBoost and 10.1 $\mathrm{m\,s^{-1}}$ for Ridge. Figure 8 illustrates this behaviour in one held-out high-wind session. The compressed tree response is consistent with sparse high-wind training data.

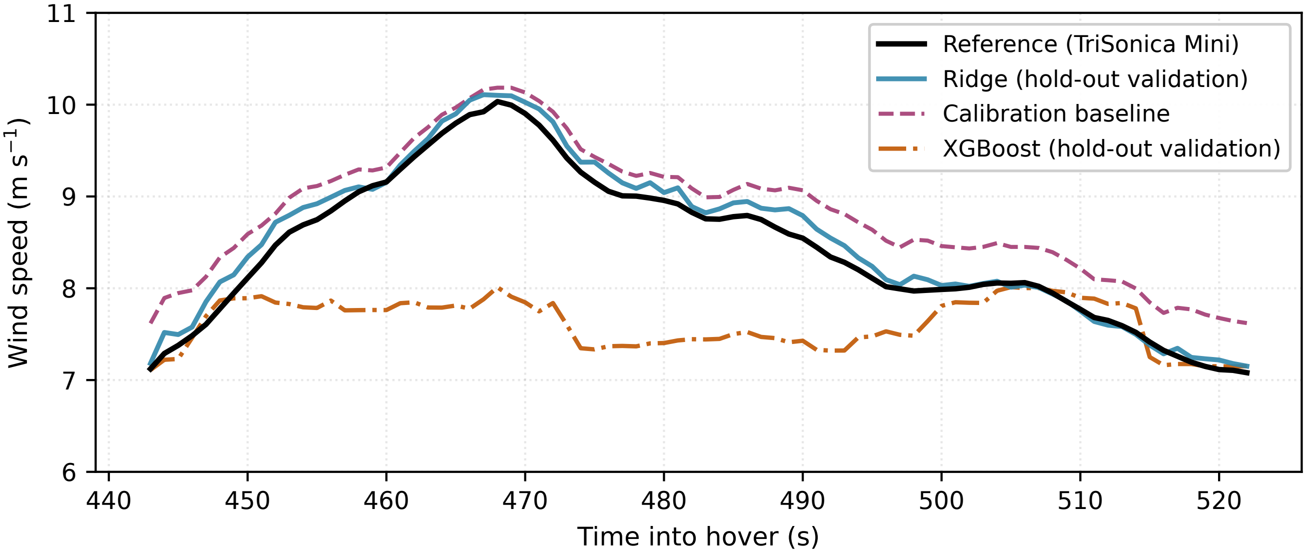

Figure 8. Wind-speed trace from a held-out hover session recorded on 18 November 2025 (session 2025-11-18 [12-28-53]). This segment has a mean reference speed of 8.4 $\mathrm{m\,s^{-1}}$, and the full-session mean is 7.6 $\mathrm{m\,s^{-1}}$. Traces show the reference anemometer, held-out Ridge and XGBoost predictions, and the calibration baseline. XGBoost varies less with the reference speed than the other estimates in this segment. The y-axis begins at 6 $\mathrm{m\,s^{-1}}$.

4 Discussion

4.1 Error characteristics and validation stability

In sessions with mean winds below 2 $\mathrm{m\,s^{-1}}$, the absolute window-mean pitch and roll are typically around 1–2° per axis, so attitude noise and station-keeping corrections have a larger relative influence. The weaker tilt signal coincides with greater scatter in the retrieved direction (Fig. 7b), which is why direction error is also reported separately for sessions above 2 $\mathrm{m\,s^{-1}}$.

The 50 sessions span 15 measurement days, and flights from the same day share meteorological conditions. Ridge has the lower mean speed RMSE on 13 of those 15 days, so its advantage is not confined to a few of them.

Across all sessions, the paired test gives $t_{49} = 5.79$ and $p < 0.001$. Resampling individual sessions yields a 95 % confidence interval of $[0.051, 0.101]$ $\mathrm{m\,s^{-1}}$ for the mean reduction in speed RMSE from the baseline to Ridge. When whole days are resampled, the interval is $[0.045, 0.108]$ $\mathrm{m\,s^{-1}}$. The same test on the 44 sessions with mean winds up to 6 $\mathrm{m\,s^{-1}}$ gives $t_{43} = 6.55$

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

and $p < 0.001$. Resampling the 14 days represented in this subset produces an interval of $[0.057, 0.120]$ $\mathrm{m\,s^{-1}}$ (Supplement Sect. S1.21).

Holding out complete measurement days instead of individual sessions changes speed RMSE by 0.009 $\mathrm{m\,s^{-1}}$ and direction MAE by less than $0.1°$ (Supplement Sect. S1.22). The confidence intervals do not include the uncertainty introduced by selecting Ridge after comparing the three trained models.

4.2 Reference limitations and physical error sources

The conservative directional uncertainty range of 5–7° includes reference alignment, aircraft heading, residual mounting effects, and the separation between the two platforms (Supplement Sect. S1.2). The observed direction MAE is comparable to this range, but retrieval error cannot be separated from reference error. The small mean signed errors above 2 $\mathrm{m\,s^{-1}}$ argue against a single constant heading offset. They do not show that the remaining error is random, because direction-dependent systematic errors can cancel in the signed mean.

Reference uncertainty also limits the interpretation of speed error. The aggregate Ridge RMSE of 0.32 $\mathrm{m\,s^{-1}}$ includes retrieval error, reference error, and microscale wind differences across the 2 m separation. In the paired analysis, Ridge reduces mean speed RMSE by 0.08 $\mathrm{m\,s^{-1}}$ relative to the baseline. This difference is smaller than the sensor's stated wind-speed accuracy of ± 0.2 $\mathrm{m\,s^{-1}}$ from 0 to 10 $\mathrm{m\,s^{-1}}$ and $\pm 2\,\%$ from 11 to 30 $\mathrm{m\,s^{-1}}$ (LI-COR, 2025). The measurements support the ranking within this setup, but they cannot attribute the full difference to retrieval performance alone.

When all six high-wind sessions are included, both methods have higher speed RMSEs than in the medium-wind regime (Sect. 3.4). After the relative-lag session is removed, the Ridge RMSE of 0.37 $\mathrm{m\,s^{-1}}$ remains above its medium-wind value of 0.29 $\mathrm{m\,s^{-1}}$. The baseline RMSE falls to 0.36 $\mathrm{m\,s^{-1}}$, below its medium-wind value of 0.38 $\mathrm{m\,s^{-1}}$. A separate gustiness analysis of all six high-wind sessions, including the relative-lag session, finds no positive relationship between Ridge absolute speed error and horizontal wind variability (Supplement Sect. S1.23). Vertical motion may also contribute, but this effect cannot be evaluated because aircraft altitude and vertical wind were not retained in the aligned analysis products.

4.3 Comparison with previous work

The Air 3S errors are of the same order as those reported for previous attitude-based retrievals, although the studies use different metrics. Across four hover flights, the indirect tilt method of Palomaki et al. (2017) gave speed RMSEs of 0.31–0.88 $\mathrm{m\,s^{-1}}$ and direction RMSEs of 10–21°. The wind-tunnel-calibrated model of Neumann and Bartholmai (2015) reached RMSEs of 0.60 $\mathrm{m\,s^{-1}}$ and 14.0° during hover with 20 s moving averages. Tilt-based wind retrieval has also been demonstrated with the purpose-built CopterSonde (Segales et al., 2020) and with commercial DJI Phantom 4 aircraft in field campaigns in Switzerland and Norway (Meier et al., 2022).

Studies using dedicated research quadrotors with per-airframe calibration against reference towers report average errors below 0.3 $\mathrm{m\,s^{-1}}$ and 8° (Wetz et al., 2021). Simma et al. (2020) obtained errors of 0.26–0.29 $\mathrm{m\,s^{-1}}$ and 4.1–4.9° over 0–5 $\mathrm{m\,s^{-1}}$. Using telemetry from 1,750 free-flight missions, García-Gascón et al. (2026) evaluated deep-learning retrievals and identified spatial variability relative to a fixed ground reference as a limitation.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

Direct comparisons are limited by differences in averaging windows, error metrics, reference setups, flight modes, and sampled wind ranges. For the Air 3S, pooling all analysis windows gives a speed RMSE of 0.35 $\mathrm{m\,s^{-1}}$ across all wind speeds and a direction MAE of 5.7° above 2 $\mathrm{m\,s^{-1}}$. These values were obtained with an unmodified consumer airframe.

4.4 Implications for field campaigns

The choice between the two methods depends on the available calibration data and the required accuracy. In operational deployment, both methods estimate wind purely from onboard telemetry without a reference anemometer. Their distinction lies in how those initial relationships are obtained: Ridge requires colocated reference-anemometer measurements during an initial training phase, whereas the calibration baseline relies on translational flights in near-calm conditions without an external reference. Ridge gives the lower aggregate speed RMSE, at 0.32 compared with 0.39 $\mathrm{m\,s^{-1}}$. In the six sessions above 6 $\mathrm{m\,s^{-1}}$, the baseline and Ridge have similar speed RMSEs of 0.48 and 0.49 $\mathrm{m\,s^{-1}}$, both lower than the errors of the tree models (Table 4). Near the upper end of the Ridge training range, the baseline therefore provides an independent cross-check based on a separate calibration. Transferring either method to another unit of the same aircraft model would require validation, followed by recalibration legs for the baseline or retraining for Ridge if unit-to-unit differences exist.

The results apply to 30 s vector means during stationary hover with one airframe at two sites. Holding out complete sessions tests transfer to other sessions within this campaign rather than to a new deployment. Because only six sessions had mean winds above 6 $\mathrm{m\,s^{-1}}$ and the maximum session mean was 7.6 $\mathrm{m\,s^{-1}}$, performance in sustained stronger winds has not been evaluated here.

4.5 Limitations

The two methods make different assumptions about airframe response. The calibration baseline treats pitch and roll as independent axes, while Ridge incorporates nonlinear and cross-axis terms fitted directly to hover data. Although the present dataset does not show elevated errors under diagonal winds, this does not establish that axis independence holds generally (Supplement Sect. S1.20). The baseline also exhibits larger crosswind-component errors than Ridge, consistent with its weaker roll calibration (Table S5). Neither method accounts for air density variations, which may influence the tilt–airspeed relationship across seasonal temperature or pressure changes.

Operational scope is also constrained by the sampling envelope. Measurements were conducted at a single nominal hover height of 1.6 m without vertical profiling, leaving boundary-layer performance at altitude unverified. All observations originate from a single airframe unit, and performance on other units or different multicopter platforms was not assessed. While the coefficients in Table 1 apply specifically to the tested DJI Air 3S unit, the calibration protocol and quadratic mapping can be repeated for other rotary-wing platforms.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

5 Conclusions

Internal telemetry from an unmodified DJI Air 3S provided horizontal wind estimates during stationary hover across 50 validation sessions. Ridge had a speed RMSE of 0.32 $\mathrm{m\,s^{-1}}$ and a direction MAE of 7.0°, compared with 0.39 $\mathrm{m\,s^{-1}}$ and 8.4° for the calibration baseline. In sessions with mean winds above 2 $\mathrm{m\,s^{-1}}$, direction MAE fell to 6.0° for Ridge and 7.0° for the baseline.

Ridge performed better across most of the sampled wind range, while the baseline reached a similar speed RMSE in the six sessions above 6 $\mathrm{m\,s^{-1}}$. During deployment, both methods estimate wind purely from onboard telemetry without an external wind sensor. Their practical trade-off is how those initial relationships are established: Ridge requires initial training against reference anemometer data, whereas the baseline derives its quadratic curves from near-calm translational flights. Neither gradient-boosted tree model improved on Ridge.

These results apply to near-surface measurements from one aircraft at two sites. Further tests are needed with additional units of the same aircraft model and with different platforms, at greater heights, and in sustained winds above the maximum session mean of 7.6 $\mathrm{m\,s^{-1}}$. Calibration flights with finer speed increments would help constrain the calibration curves for an individual airframe.

Code and data availability. The processed data and analysis code are publicly available on Zenodo at https://doi.org/10.5281/zenodo.19098691 (Polle et al., 2026). The archive contains time-aligned 10 Hz series, the 62-feature matrix, aggregated calibration measurements, fitted models, held-out predictions, validation metrics, and all processing and plotting scripts needed to reproduce the article findings. The aligned 10 Hz files (`data/processed/aligned_10hz/*_aligned.csv`) are reduced analysis products: they retain reference horizontal wind speed and direction and the derived drone-body horizontal components used in the analysis (`wind_u_body` and `wind_v_body`). They do not contain vertical wind, sonic temperature, or TriSonica diagnostic flags, and no separate fuller TriSonica export is included. Raw DJI FlightRecord files are excluded because they contain precise GNSS flight trajectories and device serial numbers. Raw TriSonica datalogger text files are likewise not included. Raw instrument logs are available from the corresponding author upon reasonable request. The reference logger software is available at https://github.com/a-polle/Trisonica_Datalogger_RasberryPi (Polle, 2026).

Author contributions. APol led the conceptualization and methodology, developed the software, conducted the measurements and formal analysis, curated the data, prepared the visualizations, and wrote the original draft. LG, APou, MV, and TW contributed to the conceptualization. LG provided scientific input throughout the project and reviewed and edited the manuscript. APou reviewed and edited successive drafts and acquired funding for the reference instrument. MV provided scientific guidance and reviewed and edited the manuscript. TW supervised the project, provided scientific input throughout, and reviewed and edited the manuscript.

Competing interests. The authors declare that they have no conflict of interest.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

330 *Acknowledgements.* This work formed part of Alexander Polle's master's thesis at the University of Bremen.

Financial support. The TriSonica anemometer was purchased with funding awarded to APou by the University of Bremen Central Research Development Fund through its Impulse grant scheme.

https://doi.org/10.5194/egusphere-2026-5393
Preprint. Discussion started: 29 September 2026

References

Bramati, M., Schön, M., Schulz, D., and Platis, A.: A Versatile Calibration Method for Rotary-Wing UAS as Wind Measurement Systems, Journal of Atmospheric and Oceanic Technology, 41, 25–43, https://doi.org/10.1175/JTECH-D-23-0010.1, 2024.

Chen, D., Su, W., Jiang, S., Yang, H., Zhang, C., Jiang, S., Chang, D., Liang, Y., Wang, H., Yang, X., Fu, T.-M., Zeng, Z., Zhu, L., Shen, H., Wang, C., and Ye, J.: Wind estimation based on flight dynamics of unmanned aerial vehicle: influencing variables and its environmental application, Atmospheric Chemistry and Physics, 26, 3607–3620, https://doi.org/10.5194/acp-26-3607-2026, 2026.

Chen, T. and Guestrin, C.: XGBoost: A Scalable Tree Boosting System, in: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16, pp. 785–794, ACM, https://doi.org/10.1145/2939672.2939785, 2016.

Crowe, D., Pamula, R., Cheung, H. Y., and De Wekker, S. F. J.: Two Supervised Machine Learning Approaches for Wind Velocity Estimation Using Multi-Rotor Copter Attitude Measurements, Sensors, 20, 5638, https://doi.org/10.3390/s20195638, 2020.

DJI: DJI Air 3S Downloads and Release Notes, https://www.dji.com/downloads/products/air-3s, last access: 28 April 2026, 2025.

European Union: Commission Implementing Regulation (EU) 2019/947 of 24 May 2019 on the rules and procedures for the operation of unmanned aircraft, http://data.europa.eu/eli/reg_impl/2019/947/oj, 2019.

García-Gascón, C., Castelló-Pedrero, P., Bas-Bolufer, J., and Chinesta, F.: Deep learning-based wind estimation from multirotor UAV telemetry, Aerospace Science and Technology, 177, 112 184, https://doi.org/10.1016/j.ast.2026.112184, 2026.

Ghirardelli, M., Kral, S. T., Cheynet, E., and Reuder, J.: SAMURAI-S: Sonic Anemometer on a MUlti-Rotor drone for Atmospheric turbulence Investigation in a Sling load configuration, Atmospheric Measurement Techniques, 18, 2103–2124, https://doi.org/10.5194/amt-18-2103-2025, 2025.

González-Rocha, J., Woolsey, C. A., Sultan, C., and De Wekker, S. F. J.: Sensing Wind from Quadrotor Motion, Journal of Guidance, Control, and Dynamics, 42, 836–852, https://doi.org/10.2514/1.G003542, 2019.

Grinsztajn, L., Oyallon, E., and Varoquaux, G.: Why do tree-based models still outperform deep learning on typical tabular data?, in: Advances in Neural Information Processing Systems, vol. 35, pp. 507–520, https://doi.org/10.52202/068431-0037, 2022.

Hattenberger, G., Bronz, M., and Condomines, J.-P.: Estimating Wind Using a Quadrotor, International Journal of Micro Air Vehicles, 14, https://doi.org/10.1177/17568293211070824, 2022.

Hoerl, A. E. and Kennard, R. W.: Ridge Regression: Biased Estimation for Nonorthogonal Problems, Technometrics, 12, 55–67, https://doi.org/10.1080/00401706.1970.10488634, 1970.

Ke, G., Meng, Q., Finley, T., Wang, T., Chen, W., Ma, W., Ye, Q., and Liu, T.-Y.: LightGBM: A Highly Efficient Gradient Boosting Decision Tree, in: Advances in Neural Information Processing Systems, vol. 30, pp. 3146–3154, Curran Associates, Inc., https://doi.org/10.5555/3294996.3295074, 2017.

LI-COR: LI-550 TriSonica Mini Wind and Weather Sensor, https://www.licor.com/products/trisonica/LI-550-mini, last access: 11 December 2025, 2025.

Meier, K., Hann, R., Skaloud, J., and Garreau, A.: Wind Estimation with Multirotor UAVs, Atmosphere, 13, 551, https://doi.org/10.3390/atmos13040551, 2022.

Neumann, P. P. and Bartholmai, M.: Real-time wind estimation on a micro unmanned aerial vehicle using its inertial measurement unit, Sensors and Actuators A: Physical, 235, 300–310, https://doi.org/10.1016/j.sna.2015.09.036, 2015.

Palomaki, R. T., Rose, N. T., van den Bossche, M., Sherman, T. J., and De Wekker, S. F. J.: Wind Estimation in the Lower Atmosphere using Multirotor Aircraft, Journal of Atmospheric and Oceanic Technology, 34, 1183–1191, https://doi.org/10.1175/JTECH-D-16-0177.1, 2017.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E.: Scikit-learn: Machine Learning in Python, Journal of Machine Learning Research, 12, 2825–2830, https://jmlr.org/papers/v12/pedregosa11a.html, 2011.

Polle, A.: Trisonica Datalogger for Raspberry Pi, https://github.com/a-polle/Trisonica_Datalogger_RasberryPi, last access: 19 March 2026, 2026.

Polle, A., Grosch, L., Poulidis, A. P., Vrekoussis, M., and Warneke, T.: Wind estimation from DJI Air 3S hover telemetry: dataset and code, Zenodo [data set], https://doi.org/10.5281/zenodo.19098691, 2026.

Reuter, M., Bovensmann, H., Buchwitz, M., Borchardt, J., Krautwurst, S., Gerilowski, K., Lindauer, M., Kubistin, D., and Burrows, J. P.: Development of a small unmanned aircraft system to derive CO_2 emissions of anthropogenic point sources, Atmospheric Measurement Techniques, 14, 153–172, https://doi.org/10.5194/amt-14-153-2021, 2021.

Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder, B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., and Dormann, C. F.: Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure, Ecography, 40, 913–929, https://doi.org/10.1111/ecog.02881, 2017.

Segales, A. R., Greene, B. R., Bell, T. M., Doyle, W., Martin, J. J., Pillar-Little, E. A., and Chilson, P. B.: The CopterSonde: an insight into the development of a smart unmanned aircraft system for atmospheric boundary layer research, Atmospheric Measurement Techniques, 13, 2833–2848, https://doi.org/10.5194/amt-13-2833-2020, 2020.

Simma, M., Mjøen, H., and Boström, T.: Measuring Wind Speed Using the Internal Stabilization System of a Quadrotor Drone, Drones, 4, 23, https://doi.org/10.3390/drones4020023, 2020.

Soltaninezhad, M., Monsorno, R., and Tondini, S.: A Review of Methods and Challenges for Wind Measurement by Small Unmanned Aerial Vehicles, Meteorological Applications, 32, e70 065, https://doi.org/10.1002/met.70065, 2025.

Thielicke, W., Hübert, W., Müller, U., Eggert, M., and Wilhelm, P.: Towards Accurate and Practical Drone-Based Wind Measurements with an Ultrasonic Anemometer, Atmospheric Measurement Techniques, 14, 1303–1318, https://doi.org/10.5194/amt-14-1303-2021, 2021.

Vauvillier, L.: dji-log-parser, https://github.com/lvauvillier/dji-log-parser, last access: 17 December 2025, 2025.

Wetz, T., Wildmann, N., and Beyrich, F.: Distributed wind measurements with multiple quadcopter unmanned aerial vehicles in the atmospheric boundary layer, Atmospheric Measurement Techniques, 14, 3795–3814, https://doi.org/10.5194/amt-14-3795-2021, 2021.

Yang, Y., Zhang, Y., Han, T., Xie, C., Liu, Y., Huang, Y., Zhou, J., Sun, H., Zhao, D., Zhang, K., and Li, S.-M.: A correction algorithm for rotor-induced airflow and flight attitude changes during three-dimensional wind speed measurements made from a rotary unoccupied aerial vehicle, Atmospheric Measurement Techniques, 18, 3035–3050, https://doi.org/10.5194/amt-18-3035-2025, 2025.

Zhu, S., Zhao, T., Zhang, H., Chen, Y., Yang, D., Liu, Y., and Cao, J.: UAVs' Flight Dynamics Is All You Need for Wind Speed and Direction Measurement in Air, Drones, 9, 466, https://doi.org/10.3390/drones9070466, 2025.

Zimmerman, S., Nagamune, R., and Rogak, S.: Wind Estimation by Multirotor Drone State Using Machine Learning with Data Rotation and Reduction, Measurement, 199, 111 491, https://doi.org/10.1016/j.measurement.2022.111491, 2022a.

Zimmerman, S., Yeremi, M., Nagamune, R., and Rogak, S.: Wind Estimation by Multirotor Dynamic State Measurement and Machine Learning Models, Measurement, 198, 111 331, https://doi.org/10.1016/j.measurement.2022.111331, 2022b.