

Sector-Conditioned Sparse Wind-Direction Sensing in Complex Terrain: Purged Field Validation, Failure Resilience and Control-Oriented Assessment

Mingyang HUANG

B2301335@ustb.edu.cn

University of Science and Technology Beijing

Article

Keywords: complex terrain, distributed sensing, wind direction, yaw control, FLORIS, fault resilience

Posted Date: October 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10877810/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Sector-Conditioned Sparse Wind-Direction Sensing in Complex Terrain: Purged Field Validation, Failure Resilience and Control-Oriented Assessment

Authors: Mingyang HUANG^{a, b, c *}

^{*}Corresponding author: Mingyang HUANG (Email: B2301335@ustb.edu.cn)

Affiliations:

^a Key Laboratory of Intelligent Bionic Unmanned Systems, Ministry of Education, University of Science and Technology Beijing, Beijing 100083, P. R. China;

^b School of Artificial Intelligence, University of Science and Technology Beijing, Beijing 100083, P. R. China;

^c Institute of Artificial Intelligence, University of Science and Technology Beijing, Beijing 100083, P. R. China;

[†]These authors contributed equally to this work

ABSTRACT

Wind-direction estimates are required for yaw control, yet sensor relevance varies across complex terrain. A four-sector, three-source topology was evaluated at 34 pre-test-eligible Perdigao sites in five purged folds. Mean target-fold RMSE decreased from 18.825 to 18.512 deg; the paired improvement was 0.313 deg (95% CI 0.085-0.570). In a fully paired 33-site cohort, improvement was 0.400 deg and remained positive under long-block and target-and-fold resampling. It exceeded nearest, kernel and circular-ridge RMSE baselines. Under causal yaw constraints, yaw RMSE decreased by 0.125 deg and the paired FLORIS farm metric improved by 0.194 percentage points. All 36 speed, turbulence and layout scenario means were positive, although two high-speed, high-turbulence intervals included zero. Dual anchors reduced short-outage degradation. Benefits were supported under unstable but not stable stratification. Results support failure-aware sparse sensing, while energy findings remain model based and height transfer was not established.

Keywords: complex terrain; distributed sensing; wind direction; yaw control; FLORIS; fault resilience

Nomenclature.

Symbol	Definition	Unit
$\theta, \hat{\theta}$	Target and estimated wind direction	deg
$e\theta$	Wrapped angular error	deg
w_i, r_i	Source weight and target-source distance	-, m
L_d	Calibration-selected distance scale	m
ψ_k, c_k	Nacelle heading and commanded yaw target	deg
γ	Yaw misalignment	deg
η_{iso}	Isolated-turbine fractional loss proxy	-
u^*, L	Friction velocity and Obukhov length	m s ⁻¹ , m
D, TI	Rotor diameter and turbulence intensity	m, -

1. Introduction

Wind direction is a control input whose spatial and temporal variability affects yaw alignment, wake steering and structural loading. The relevant dynamics extend from seconds to minutes and from individual rotors to the wind-farm scale. Measurement uncertainty and local flow distortion can therefore be propagated directly into

controller decisions [1]. Persistent and dynamic yaw errors also remain difficult to diagnose from turbine signals alone [2].

Yaw and wake-steering controllers are commonly designed with simplified representations of inflow direction. Their performance is reduced when direction variability is ignored or treated as stationary [3]. Utility-scale demonstrations have shown that model-data integration can improve controller deployment, although accurate ambient-direction estimates remain a prerequisite [4]. Forecast errors can also alter the optimum yaw settings and their expected power benefit [5].

In complex terrain, the sensing problem becomes spatial rather than purely temporal. Ridges, valleys, roughness changes and thermal circulations can turn or decouple the flow over short distances. Wind turbines and meteorological towers can consequently be interpreted as a distributed sensor array. Operational data have already been used to infer orographic and plant-induced heterogeneity [6], while Gaussian-process models have been applied to ambient wind-field estimation [7].

Dense estimation is not necessarily equivalent to informative estimation. Additional sensors can be redundant, unavailable or exposed to flow regimes that do not represent a local target. Field measurements and numerical fields have therefore been fused to reconstruct winds in mountainous terrain [8]. However, the value of an individual measurement remains conditional on direction, stability and topology. A sparse and interpretable selection rule may be preferable when communication, maintenance or data quality limits the available network.

The Perdigao campaign provides a demanding basis for this question. The double-ridge site exhibits separation, recirculation, channelled flow and pronounced spatial variability [9,10]. Recent multiscale simulations have further shown that higher resolution improves the representation of terrain-driven structures without guaranteeing uniformly lower pointwise errors [11]. These characteristics make a leakage-controlled field test more informative than an unconstrained parameter sweep.

This study tests whether a fixed four-sector sparse topology improves unseen wind-direction estimation relative to a static three-source topology. The estimator is then passed through a causal yaw controller and a standardized FLORIS layout. A dual-anchor extension is evaluated under simulated sensor outages. Atmospheric stability, turbulence intensity and terrain descriptors are used to identify performance boundaries. Five contributions are provided: purged temporal validation, comparison with strong equal-budget baselines, target-wise robustness, multi-condition control translation, and failure-aware regime analysis.

2. Field data and validation design

2.1 Perdigao observations


The primary observations were obtained from the NSF NCAR/EOL quality-controlled five-minute Integrated Surface Flux System product, version 1.2 [12]. Forty-five daily files from 1 May to 14 June 2017 were processed. Wind direction, speed, horizontal velocity variances and sonic covariance variables were retained. Coordinates and surveyed tower altitudes were extracted from the same archive.

Direction was accepted within $[0, 360)$ deg, and wind speed was restricted to $2\text{--}30\text{ m s}^{-1}$. A record was rejected when the logical diagnostic was non-zero. Turbulence intensity was calculated from the horizontal velocity variances and was retained when finite and not greater than 3. The 10 m instrument at each target was treated as a local reference rather than an error-free representation of rotor-equivalent inflow.

Site eligibility was fixed before outer testing by requiring at least 500 valid direction records before 21 May 2017. This calibration-only rule retained 34 sites and did not inspect test-period direction, error or availability when defining the primary population. A fully paired downstream cohort of 33 sites was formed only for analyses that required estimable metrics in every outer fold. The excluded site and the reason for its exclusion were fixed before the control and FLORIS translations. Results from the primary population and the evaluation-complete cohort are reported separately.

Terrain and 10 m tower network

Perdigão, Portugal



Source: Copernicus DEM GLO-30 Public; NCAR/EOL Perdigao ISFS tower metadata.

Figure 1. Terrain context and the NCAR 10 m tower network at Perdigao.

2.2 Purged temporal validation

Five non-overlapping outer tests were started on 21 May, 26 May, 31 May, 5 June and 10 June. Each test covered five days. The preceding 14 days formed the outer calibration window, which ended 12 h before the test. Topology, distance scale and estimator parameters were fitted only in the purged calibration interval. Each target was excluded from its own source set.

A fixed four-sector design was treated as the primary engineering configuration. An additional nested audit compared zero, four, eight and twelve sectors. Three rolling two-day inner validation blocks were placed within each outer calibration interval. A 6 h gap separated every inner training and validation segment. The selected sector count was then refitted on the full outer calibration interval and evaluated once on the unseen outer test.

Validation workflow

Field estimates are validated before control and model-based energy assessment.




Figure 2. Leakage-controlled validation and control-assessment workflow.

Table 1. Data, validation roles and permitted interpretation.

Evidence layer	Data and scale	Validation role	Permitted interpretation
Primary network	34 pre-test-eligible sites; 10 m; 5 min	Five purged 5-day outer tests	Near-surface direction estimation
Downstream cohort	33 evaluation-complete sites	Paired baselines, control and FLORIS	Control-relevant model assessment

Evidence layer	Data and scale	Validation role	Permitted interpretation
Height audit	Three common towers; 10 and 100 m	Purged transfer test	Boundary on height transfer
Control translation	Continuous test records	Causal yaw dynamics	Tracking and modelled energy
Sensitivity audit	Balanced target-fold sample	36 FLORIS and 6 controller cases	Dependence on model settings
Terrain	Copernicus GLO-30	Explanatory descriptors	Heterogeneity, not flow truth

3. Methods

3.1 Circular estimation and sparse topology

Angular errors were evaluated on the circle. For an estimate $\hat{\theta}$ and a target direction θ , the wrapped error was defined as

$$e_{\theta} = [(\hat{\theta} - \theta + 180^{\circ}) \bmod 360^{\circ}] - 180^{\circ} \quad (1)$$

A fused direction was calculated from the finite source directions θ_i and non-negative weights w_i :

$$\hat{\theta} = \text{atan2}(\sum_i w_i \sin \theta_i, \sum_i w_i \cos \theta_i) \bmod 360^{\circ} \quad (2)$$

Distance weighting was applied after topology selection. The source weight was

$$w_i = \exp(-\frac{r_i}{L_d}) \quad (3)$$

where r_i is the target-source distance and L_d is the calibration-selected distance scale. Candidate scales of 250, 500, 1000, 2000 and 4000 m were tested without using the outer evaluation period.

The static topology was selected greedily with a three-source runtime budget. At each step, the candidate that minimized calibration RMSE after circular fusion was added. For a selected set S_j , the next source was

$$s_{j+1} = \underset{s \notin S_j}{\text{argmin}} \text{RMSE}_{\text{cal}}(S_j \cup \{s\}) \quad (4)$$

The first static source was designated as the anchor. Its observed direction defined four 90 deg runtime sectors. Within each supported calibration sector, the anchor was retained and two additional sources were selected with the same greedy objective. A sector required at least 20 valid calibration target records. Unsupported sectors and missing anchor observations were assigned to the static topology. The target direction was never used as a runtime gate.

Algorithm 1. Leakage-controlled sector-conditioned sparse estimation.

Step	Calibration or runtime operation	Leakage control
1	Freeze target, candidate nodes and purged calibration indices	Outer-test records remain inaccessible
2	Tune distance scale and select static greedy-3 topology	Calibration RMSE only
3	Use the first static node as the directional anchor	No target direction enters the gate
4	Fit two additional sources in each supported 90 deg sector	At least 20 calibration target records
5	At runtime, gate by anchor direction and fuse three finite sources	Static fallback for unsupported or missing gates
6	Score the frozen estimate once on the outer test	No post-test tuning

For K active sources, M candidates and N calibration records, exhaustive candidate scoring within the greedy sequence scales as $O(K^2MN)$ per topology. Four sector fits multiply this calibration cost by at most four. Runtime circular fusion is $O(K)$ per record and used $K = 3$ throughout. Thus the added computation occurs primarily offline; online cost remains constant with respect to network size after topology freezing.

3.2 Comparator models

Six comparators were evaluated. Nearest-three, static greedy-three, sector-conditioned greedy-three and sparse circular ridge shared an equal three-node runtime budget. The directional kernel and dense circular ridge used all candidate sources. Ridge features comprised sine, cosine and availability indicators, with calibration-mean imputation for missing values.

Ridge penalties of 0.01, 0.1, 1, 10 and 100 were evaluated in three rolling two-day inner blocks separated by 6 h gaps. RMSE was primary; mean absolute error and P95 absolute error were retained to distinguish central accuracy from suppression of larger errors.

Table 2. Frozen estimator definitions and runtime budgets.

Method	Runtime sources	Calibration operation	Runtime fallback
Static greedy-3	3	Target-specific greedy sequence	Available selected sources
Sector-4 greedy-3	3 per sector	Four anchor-direction sequences	Static greedy-3
Nearest-3	3	Distance ranking	Available nearest sources
Sparse ridge-3	3	Inner source and penalty tuning	Mean-imputed features
Dense ridge	All 32 candidates	Inner penalty tuning	Mean-imputed features
Kernel	All 32 candidates	Directional kernel	Available candidates

3.3 Causal yaw controller and FLORIS translation

The outer-test estimates were passed to a variable-time-step yaw controller. At each record, nacelle movement was first calculated under the previous command. Tracking error was then scored, and the current estimate was allowed to update only future commands. This order prevented a current observation from changing the nacelle state at the same instant.

Yaw rate was limited to 0.30 deg s⁻¹. A new setpoint required a 2 deg deadband and a minimum dwell of 20 s. A 600 s heartbeat refreshed a persistent command. When an estimate was missing, the previous yaw target was held. The resulting heading ψ_k was advanced using

$$\psi_k = \psi_{k-1} + \text{clip}[\text{wrap}_{180}(c_{k-1} - \psi_{k-1}), -0.30\Delta t_k, 0.30\Delta t_k] \quad (5)$$

where c_{k-1} is the previous command and Δt_k is the elapsed time in seconds. Commands per hour, setpoint changes per hour and yaw travel per hour were retained as operational-cost indicators.

Energy consequences were evaluated with FLORIS version 4 and the Gauss-Curl Hybrid model [13]. A nine-turbine NREL 5 MW layout was arranged as a rotated 3 x 3 array with 5D streamwise and 3D crosswind spacing, where $D = 125.88$ m. Wind speed was standardized at 8 m s⁻¹, while observed turbulence intensity was clipped to 0.04-0.25. FLORIS has been used extensively for control-oriented wake evaluation and utility-scale control design [14].

The isolated-turbine loss proxy was calculated from yaw misalignment γ as

$$\eta_{\text{iso}} = 1 - \max[\cos(\gamma), 0]^{1.88} \quad (6)$$

Equation (6) was used only as a transparent control-oriented translation. It was not interpreted as measured electrical production.

The isolated-turbine retained-power fraction was calculated as $\max(\cos \gamma, 0)^{1.88}$, where γ is yaw misalignment. Farm power was computed only for paired records with absolute yaw misalignment not greater than 45 deg. Both methods were evaluated on identical records. These calculations quantify model-based consequences and do not represent measured annual energy production.

3.4 Dual-anchor resilience

A second anchor was added to reduce dependence on the primary gate sensor. The first and second static sources were assigned as primary and backup anchors. A circular calibration-only offset corrected the backup gate. Each supported sector still used three active sources, and source reuse reduced the mean number of unique deployed nodes.

Permanent primary-anchor loss and bursts of 0.5, 2 and 6 h were simulated. Burst onsets were placed at 10%, 30%, 50%, 70% and 90% of each outer test. Outage-only RMSE required at least three paired records. The experiments represent data-channel failures rather than physical turbine component faults. This distinction is consistent with the broader separation between fault diagnosis, redundancy and fault-tolerant control [15].

3.5 Stability, turbulence and terrain

Atmospheric stability was derived from measured 10 m sonic covariances. The friction velocity and Monin-Obukhov length proxy were calculated as

$$u_* = \sqrt[4]{(u'w')^2 + (v'w')^2} \quad (7)$$

$$L = -\frac{T_s u_*^3}{\kappa g w' T_s'} \quad (8)$$

where T_s is sonic temperature in kelvin, $\kappa = 0.4$ and $g = 9.80665 \text{ m s}^{-2}$. Records were retained for $0.05 \leq u_* \leq 3 \text{ m s}^{-1}$ and $|z/L| \leq 10$. Unstable, near-neutral and stable classes were defined by $z/L < -0.1$, $|z/L| \leq 0.1$ and $z/L > 0.1$. Sonic-temperature flux was used for regime stratification rather than as a precision heat-flux product. Stability classification over complex terrain is known to remain sensitive to measurement location and formulation [16].

Turbulence-intensity boundaries were fixed at 0.10 and 0.20. Wind-speed groups of 2-5, 5-10 and above 10 m s^{-1} were exploratory. Copernicus GLO-30 profiles supplied elevation difference, local slope, residual roughness and ridge-crossing descriptors. DEM descriptors were used to interpret spatial heterogeneity and were not treated as terrain-resolving flow calculations.

3.6 Metrics and uncertainty

The primary estimand was the mean of 169 valid target-fold RMSE differences in the 34-site calibration-only population, defined as static minus sector-conditioned RMSE. Positive values favour the sector design. Fold-cluster intervals were obtained by resampling the five outer folds. A secondary evaluation-complete cohort contained 165 paired target-folds from 33 sites. Temporal dependence in this cohort was examined with hierarchical 12, 24 and 48 h block resampling, with 5000 repetitions at each block length.

Regime intervals were obtained by resampling outer folds and then target clusters within each selected fold. Continuous-control differences were paired by target-fold. A result was described as supported only when its 95% interval excluded zero. Subgroup analyses were interpreted as mechanistic boundaries, and unsupported results were retained.

3.7 Spatial and control-sensitivity audits

Spatial robustness was assessed in the evaluation-complete cohort from 165 paired target-fold effects. Mean effects were summarized for each target and terrain group. A two-way bootstrap independently resampled the 33 targets and five outer folds. Each target was then removed in turn to test whether a single location dominated the network mean. This audit quantifies heterogeneity after target-specific calibration; it is not a zero-calibration transfer test for a new site.

Controller sensitivity covered yaw-rate limits of 0.15, 0.30 and 0.60 deg s^{-1} , deadbands of 1, 2 and 4 deg, and dwell settings of 20 and 60 s in six prespecified combinations. FLORIS sensitivity used ten randomly selected conditions per target-fold before yaw-domain filtering. Wind speeds of 6, 8, 10 and 12 m s^{-1} , turbulence intensities of 0.06, 0.10 and 0.18, and compact, nominal and open 3 x 3 layouts were crossed to form 36 scenarios. The balanced sample contained 1487 paired field conditions after filtering and was used for sensitivity, not annual-energy estimation.

4. Results

4.1 Purged field validation

The calibration-only primary population retained 34 sites and 169 valid target-folds. Mean static and sector-conditioned RMSE values were 18.825 and 18.512 deg. The paired reduction was 0.313 deg, equivalent to 1.7% of the static RMSE, with a fold-cluster 95% interval of 0.085-0.570 deg. Four of five fold means were positive. The positive network mean and interval support the primary effect, whereas the deteriorating fold limits any claim of universal temporal improvement.

In the evaluation-complete 33-site cohort, static and sector-conditioned RMSE values were 19.114 and 18.714 deg. The paired reduction was 0.400 deg, equivalent to 2.1%, with a fold-cluster interval of 0.233-0.550 deg. Fold improvements were 0.093, 0.443, 0.480, 0.345 and 0.641 deg. This cohort was retained for equal-sample baseline, spatial, controller and FLORIS comparisons, not used to redefine primary eligibility.

Temporal block resampling in the evaluation-complete cohort produced mean reductions of 0.393, 0.397 and 0.395 deg for 12, 24 and 48 h blocks. Their lower 95% limits were 0.078, 0.089 and 0.085 deg. The effect therefore remained positive when serial dependence was represented at timescales much longer than the five-minute sampling interval.




Figure 3. Purged field-validation results. (a) Fold-wise transfer. (b) Block-resampled uncertainty.

4.2 Strong baseline comparison

The sector-conditioned estimator had the lowest mean RMSE and P95 error among the evaluated methods. Its RMSE advantage over sparse circular ridge was 1.136 deg, with a 95% interval of 0.688-1.463 deg. Advantages over the dense ridge, nearest-three and all-node kernel methods were 1.434, 3.900 and 7.883 deg. Every corresponding fold-cluster interval remained above zero.

Metric rankings were not identical. Sparse ridge achieved a lower mean absolute error of 12.467 deg, compared with 12.854 deg for sector conditioning. Its RMSE and P95 errors were nevertheless higher. Sector conditioning was therefore associated with stronger suppression of larger angular errors rather than uniform dominance across all loss functions.




Figure 4. Estimator performance under purged field validation.

Table 3. Purged field performance of the estimator baselines.

Method	RMSE (deg)	MAE (deg)	P95 (deg)	Coverage	Active / deployed nodes
Sector-4 greedy-3	18.714	12.854	36.341	0.927	3 / 5.84
Static greedy-3	19.114	13.294	36.724	0.927	3 / 3
Sparse ridge-3	19.849	12.467	38.179	0.927	3 / 3
Dense ridge	20.148	12.708	40.129	0.999	32 / 32
Nearest-3	22.614	16.363	43.866	0.937	3 / 3
All-node kernel	26.596	19.560	52.983	0.999	32 / 32

4.3 Causal yaw and modelled energy consequences

The estimator effect persisted after causal controller dynamics were introduced. Mean yaw-tracking RMSE was reduced from 32.284 to 32.159 deg. The paired reduction was 0.125 deg, with a 95% interval of 0.047-0.202 deg.

The smaller effect relative to direct estimation was expected because yaw-rate and dwell constraints filtered rapid changes.

Mean isolated-turbine energy loss decreased from 14.523% to 14.228%. The paired improvement was 0.295 percentage points, with a 95% interval of 0.243-0.349. The farm-energy change remained negative for both methods relative to the aligned reference. It changed from -5.998% to -5.803%, which corresponds to a 0.194 percentage-point paired advantage with an interval of 0.123-0.268.

Operational activity changed only slightly. Commands increased by 0.011 h⁻¹, and setpoint changes increased by 0.026 h⁻¹. Both intervals were above zero. Yaw travel increased by 0.325 deg h⁻¹, but its interval of -0.252 to 0.902 crossed zero. Neither isolated nor farm CVaR improvement was supported because both intervals included zero.




Figure 5. Causal yaw-control and FLORIS consequences.

4.4 Dual-anchor failure resilience

Normal-operation RMSE was 18.714 deg for the single-anchor design and 18.607 deg for the dual-anchor design. The 0.107 deg difference had an interval of -0.141 to 0.446 deg and was not supported. Mean unique deployed nodes were reduced from 5.84 to 4.93 because the two anchors were reused across sector-specific source sets.

Under permanent primary-anchor loss, RMSE increased to 22.700 deg for the single-anchor method and 22.010 deg for the dual-anchor method. The paired benefit was 0.691 deg, with a 95% interval of 0.387-0.874 deg. During outage windows, benefits were 1.030 deg for 0.5 h bursts and 0.805 deg for 2 h bursts. Both intervals excluded zero.

The 6 h outage-only benefit was 0.350 deg, with an interval of -0.162 to 0.819 deg. Whole-test burst differences also had intervals crossing zero because the outage occupied a limited fraction of each five-day evaluation. The evidence therefore supports short-to-moderate simulated outage resilience, not duration-independent protection.

4.5 Atmospheric and terrain boundaries

Quality-controlled stability features were available for 88.24% of the complete 33-site archive and for 98.52% of the paired estimator records. Under unstable stratification, the sector-conditioned RMSE reduction was 0.372 deg. Its hierarchical 95% interval was 0.035-0.701 deg. The near-neutral reduction was 0.303 deg, but its interval of -0.141 to 0.887 deg included zero.

Stable conditions produced a reduction of -0.013 deg, with an interval of -0.171 to 0.122 deg. No stable-flow improvement was established. The high-TI group showed a larger point estimate of 0.326 deg, although its lower interval limit was -0.013 deg. In an exploratory speed split, the 5-10 m s⁻¹ group improved by 0.541 deg with an interval of 0.172-1.043 deg.

DEM regularization changed only 18.8% of target-fold topologies in the separate terrain audit. It reduced mean source distance by 5.8% and mean path ridge excess by 11.2%. Its first outer fold did not improve, and target-level inner selection showed weak external transfer. Terrain descriptors were therefore retained as explanatory evidence rather than as a replacement estimator.




Figure 6. Effects of anchor outages and atmospheric stability.

4.6 Spatial and control robustness

Twenty-four of 33 targets had a positive five-fold mean effect. The median target effect was 0.071 deg, with an interquartile range of -0.003 to 0.463 deg, which confirms substantial spatial heterogeneity. The two-way target-and-fold bootstrap interval was 0.027-0.873 deg. Removing any single target left a positive network mean between 0.274 and 0.443 deg.

Terrain groups were not equally responsive. The northwest transect showed the largest mean improvement of 0.942 deg. Mean effects were 0.400 deg in the valley and 0.384 deg on the southeast transect, whereas the northeast and southwest ridges improved by only 0.032 and 0.035 deg. The method should therefore be interpreted as a network-level design with location-dependent value.

Across the six controller configurations, yaw-RMSE improvements ranged from 0.087 to 0.143 deg and every fold-cluster interval excluded zero. Isolated-loss improvements ranged from 0.256 to 0.297 percentage points. The 20 and 60 s dwell cases were identical because the observations were sampled every 5 min; dwell sensitivity below that interval cannot be identified from these data.

All 36 FLORIS scenario means favoured sector conditioning, with paired advantages from 0.120 to 0.350 percentage points. Thirty-four fold-cluster intervals excluded zero. The two unsupported cases occurred at 12 m s⁻¹ and turbulence intensity 0.18 in the nominal and open layouts. The benefit generally weakened as wind speed and turbulence intensity increased.




Figure 7. Target-wise effects and FLORIS scenario sensitivity.

4.7 Negative validation audits

Nested per-target selection among zero, four, eight and twelve sectors did not transfer reliably. Its mean reduction was 0.251 deg, but the 95% interval was -0.132 to 0.531 deg. The fixed four-sector configuration was therefore retained as the tested engineering design. It was not interpreted as the outcome of a validated adaptive sector-count controller.

Height transfer was examined at tse04, tse09 and tse13, where 10 and 100 m observations were jointly available. The 100 m reduction was -0.052 deg, with an interval of -0.301 to 0.197 deg. The available three-tower analysis did not establish vertical transfer. The primary result consequently remains a near-surface network result.

Table 4. Paired effects, confidence intervals and claim decisions.

Outcome	Effect	95% interval	Decision
Primary estimator RMSE (34 sites)	+0.313 deg	0.085 to 0.570	Supported; 4/5 folds
Evaluation-complete cohort (33 sites)	+0.400 deg	0.233 to 0.550	Supported; all folds
Two-way spatial-temporal bootstrap	+0.400 deg	0.027 to 0.873	Supported; heterogeneous
Yaw-tracking RMSE	+0.125 deg	0.047 to 0.202	Supported
FLORIS farm metric	+0.194 pp	0.123 to 0.268	Supported; model based
FLORIS 36-scenario sensitivity	+0.120 to +0.350 pp	34/36 intervals positive	Broad, not universal
Permanent anchor loss	+0.691 deg	0.387 to 0.874	Supported
Unstable stratification	+0.372 deg	0.035 to 0.701	Supported
Stable stratification	-0.013 deg	-0.171 to 0.122	Not supported
Adaptive sector count	+0.251 deg	-0.132 to 0.531	Not supported
100 m transfer	-0.052 deg	-0.301 to 0.197	Not supported

5. Discussion

5.1 Scientific interpretation

The central result is that source relevance changes with the directional state of a complex-terrain network. A fixed three-source set cannot represent every flow sector equally well. Four anchor-defined sectors allowed the source set to change while preserving a constant active-node budget. The resulting effect was modest in degrees but remained positive in the calibration-only population, the evaluation-complete cohort, long-block resampling and the two-way spatial-temporal bootstrap.

The target analysis prevents the network mean from being mistaken for uniform local improvement. Most targets improved, but the median effect was much smaller than the mean and the northwest transect contributed disproportionately. No single target controlled the conclusion. This pattern suggests that conditional topology is most valuable where directional coupling changes strongly with flow sector, while fixed topologies remain adequate at weakly responsive ridge locations.

The strongest sparse comparator exposed an important distinction. Circular ridge reduced average absolute error more effectively, whereas sector conditioning reduced RMSE and P95 error. Large angular errors carry disproportionate consequences for yaw commands and wake-model inputs. A topology that suppresses these events can therefore remain operationally useful even when its mean absolute error is not minimal.

The unstable-stratification result provides a plausible physical boundary. Convective mixing, local turning and spatial variability increase the chance that a fixed source set becomes unrepresentative. Sector conditioning was more beneficial in that regime. Stable flow did not show the same effect. This pattern is consistent with the known dependence of direction variability on stability and terrain [1,16], but it does not establish a universal causal mechanism.

5.2 Estimation and control must be separated

A lower direct estimator RMSE was not assumed to produce an equivalent energy improvement. The continuous controller reduced the direct effect through actuator-rate, dwell and deadband constraints. This attenuation is desirable because it prevents rapid estimator fluctuations from being converted into immediate nacelle motion. Dynamic wake models similarly emphasize the need to represent control transients rather than rely only on steady states [17].

The FLORIS farm metric remained negative relative to perfect alignment for both estimators. Sector conditioning made the value less negative; it did not create a measured production gain. Field wake-steering studies have shown that wind-direction uncertainty, turbine response and filtering can materially change predicted gains [3,18]. A defensible energy statement must therefore remain conditional on the selected turbine, layout and model assumptions.

The multi-condition audit reduces dependence on the nominal 8 m s⁻¹ case. Positive mean differences were obtained across every tested speed, turbulence and layout combination, although two high-speed, high-turbulence intervals included zero. The control implication is therefore robust over the tested design space but weaker under

conditions in which wake recovery and power sensitivity reduce the incremental value of small direction improvements.

5.3 Reliability and deployment implications

The dual-anchor result indicates that redundancy is most valuable during a failed gate signal. Normal-operation accuracy was preserved within uncertainty, while permanent and short outage degradation was reduced. The effect became inconclusive for 6 h outages. Long disruptions likely require reconfiguration beyond a corrected backup anchor, such as re-estimation of sector topology or an independent inflow source.

A practical deployment sequence can be derived. Candidate nodes should first be screened for pre-deployment availability and calibrated against local targets. Static and sector-conditioned sparse sets should then be compared with equal-budget statistical baselines. A backup gate should be offset-corrected on calibration data. Every missing-data policy should be evaluated on error, coverage, freshness and command activity rather than availability alone.

The study does not replace turbine-level fault diagnosis or SCADA condition monitoring. Those tasks involve different sensors, failure modes and maintenance objectives [19]. Instead, the results concern resilience of a distributed meteorological information pathway that supplies a yaw-control calculation.

5.4 Scope of the contribution

The novelty lies in the complete evidence chain rather than in a new yaw actuator. A sparse sector-conditioned topology was tested with purged field folds, autocorrelation-aware intervals, equal-budget baselines, causal controller dynamics and explicit failure scenarios. Negative transfer and subgroup results were retained. This structure limits performance claims while improving reproducibility.

Four sectors and three active sources should not be treated as universal design constants. Their value depends on spatial density, terrain, instrumentation, sampling interval and control timescale. The general design principle is conditional information selection under a fixed resource budget, followed by downstream validation in the actual control logic.

6. Limitations

The primary network uses 10 m observations rather than rotor-equivalent or hub-height inflow. Only three towers provided a common 100 m audit, and no height-transfer benefit was established. Results from the near-surface network cannot therefore be assigned directly to utility-scale hub-height sensing.

The target instrument is a local measurement with calibration and flow-distortion uncertainty. It is not an error-free atmospheric truth. Eligibility for the 34-site primary population was fixed using pre-test records only, but one primary fold deteriorated. The 33-site cohort required complete downstream pairing and produced the stronger all-fold result. Reporting both populations exposes rather than conceals this dependence.

One field campaign was used. Temporal folds and target resampling quantify within-campaign robustness but do not establish zero-calibration transfer to another site, season or sensor technology. Terrain descriptors came from a 30 m DEM and cannot resolve instantaneous separation, recirculation or thermal flows. The sonic-temperature-flux stability proxy was suitable for stratification but was not a precision flux product.

FLORIS used standardized NREL 5 MW layouts and a balanced subset of field directions rather than a site-specific turbine controller, measured SCADA power or annual wind climate. Energy and CVaR quantities must remain model-based. Loads were not evaluated, yaw travel uncertainty remained inconclusive, and sub-five-minute dwell effects could not be identified from five-minute observations.

Adaptive sector-count selection, stable-flow improvement, 6 h outage resilience and 100 m transfer were not supported. These outcomes constrain the method's operational envelope. Independent hub-height networks and synchronized SCADA power would provide the most valuable next validation.

7. Conclusions

A fixed four-sector sparse topology was evaluated for distributed wind-direction estimation in complex terrain. Across 34 pre-test-eligible Perdigao 10 m sites and five purged outer tests, mean target-fold RMSE was reduced by 0.313 deg relative to a static three-source topology, with a 95% interval of 0.085-0.570 deg. In the evaluation-complete 33-site cohort, the reduction was 0.400 deg and remained positive under 12-48 h block resampling. Strong nearest, kernel and circular-ridge baselines were also exceeded on RMSE and P95 error.

The estimator effect persisted after causal yaw dynamics. Yaw RMSE decreased by 0.125 deg, isolated modelled loss decreased by 0.295 percentage points, and the paired FLORIS farm metric improved by 0.194 percentage points. All 36 FLORIS scenario means were positive, while 34 intervals excluded zero. Additional command activity was small. A dual-anchor design reduced degradation during permanent and 0.5-2 h simulated outages.

Target-wise evidence was heterogeneous but not driven by one location. Twenty-four targets had positive mean effects, the two-way bootstrap remained above zero, and every leave-one-target-out network mean was positive. These results support bounded network-level robustness rather than uniform site-level superiority.

Performance was supported under unstable stratification but not under stable conditions. Automatic sector-count selection and 100 m transfer were inconclusive. The contribution is therefore a bounded, failure-aware framework for sector-conditioned sparse sensing and downstream control assessment. Independent hub-height and SCADA power validation is required before farm-wide production claims can be made.

Acknowledgements

The Perdigao campaign participants and the NSF NCAR/EOL Integrated Surface Flux System Team are acknowledged for providing the quality-controlled field observations. Copernicus DEM GLO-30 was accessed through the public AWS Registry of Open Data.

Author contributions

Author contribution roles are omitted from the anonymized manuscript and will be supplied in the title-page file using the CRediT taxonomy.

Statements and Declarations

Ethical considerations

Ethics approval was not required because numerical simulations and publicly available non-human engineering datasets were used.

Consent to participate

Not applicable.

Consent for publication

Not applicable.

Declaration of conflicting interest

The authors declared no potential conflicts of interest with respect to the research, authorship and publication of this article.

Conflict-of-Interest Statement

The authors of this paper declare that this paper has no Conflict-of-Interest. The authors of this paper state that this paper has no Conflict-of-Interest. The authors do not have any conflicts of interest with other entities or researchers.

Acknowledgment and funding

This work was supported by the project supported by no Fund.

Funding statement

No specific funding was reported for this research.

Data availability

The primary Perdigao dataset is publicly available at <https://doi.org/10.26023/ZDMJ-D1TY-FG14>. Processing scripts, frozen configurations and scenario-level result tables have been prepared as an anonymized reproducibility package. A permanent repository URL or DOI will be added to the submission record before publication.

References

1. Dallas S, Stock A and Hart E (2024) Control-oriented modelling of wind direction variability. *Wind Energy Science* 9: 841-867. DOI: 10.5194/wes-9-841-2024.
2. Li Q, Chen D, Lin H and Yang X (2025) A review of diagnostic methods for yaw errors in horizontal axis wind turbines. *Energies* 18(3): 588. DOI: 10.3390/en18030588.
3. Simley E, Fleming P and King J (2020) Design and analysis of a wake steering controller with wind direction variability. *Wind Energy Science* 5: 451-468. DOI: 10.5194/wes-5-451-2020.
4. Bachant P, Ireland P, Burrows B et al. (2024) Development and validation of a hybrid data-driven model-based wake steering controller and its application at a utility-scale wind plant. *Wind Energy Science* 9: 2235-2259. DOI: 10.5194/wes-9-2235-2024.
5. Kadoche E, Bianchi P, Carton F et al. (2024) On the importance of wind predictions in wake steering optimization. *Wind Energy Science* 9: 1577-1594. DOI: 10.5194/wes-9-1577-2024.

6. Braunbehrens R, Vad A and Bottasso CL (2023) The wind farm as a sensor: learning and explaining orographic and plant-induced flow heterogeneities from operational data. *Wind Energy Science* 8: 691-723. DOI: 10.5194/wes-8-691-2023.
7. Van der Hoek DC, Sinner M, Simley E et al. (2021) Estimation of the ambient wind field from wind turbine measurements using Gaussian process regression. In: 2021 American Control Conference, New Orleans, LA, 25-28 May 2021, pp. 558-563. DOI: 10.23919/ACC50511.2021.9483088.
8. Cheng X, Yan B, Zhou X et al. (2024) Wind resource assessment at mountainous wind farm: fusion of RANS and vertical multi-point on-site measured wind field data. *Applied Energy* 363: 123116. DOI: 10.1016/j.apenergy.2024.123116.
9. Fernando HJS, Mann J, Palma JMLM et al. (2019) The Perdigao: peering into microscale details of mountain winds. *Bulletin of the American Meteorological Society* 100(5): 799-820. DOI: 10.1175/BAMS-D-17-0227.1.
10. Menke R, Vasiljevic N, Mann J et al. (2019) Characterization of flow recirculation zones at the Perdigao site using multi-lidar measurements. *Atmospheric Chemistry and Physics* 19: 2713-2723. DOI: 10.5194/acp-19-2713-2019.
11. Al Oqaily D, Giani P and Crippa P (2025) Evaluating WRF multiscale wind simulations in complex terrain: insights from the Perdigao field campaign. *Journal of Geophysical Research: Atmospheres* 130: e2025JD044055. DOI: 10.1029/2025JD044055.
12. NSF NCAR/EOL Integrated Surface Flux System Team (2021) NCAR/EOL quality controlled 5-minute ISFS surface flux data, geographic coordinate, tilt corrected, version 1.2 [Dataset]. DOI: 10.26023/ZDMJ-DITY-FG14.
13. National Renewable Energy Laboratory (2024) FLORIS, version 4 [Computer software]. Available at: <https://github.com/NREL/floris> (accessed 27 August 2026).
14. Howland MF, Quesada JB, Martinez JJP et al. (2022) Collective wind farm operation based on a predictive model increases utility-scale energy production. *Nature Energy* 7: 818-827. DOI: 10.1038/s41560-022-01085-z.
15. Fekih A, Habibi H and Simani S (2022) Fault diagnosis and fault tolerant control of wind turbines: an overview. *Energies* 15(19): 7186. DOI: 10.3390/en15197186.
16. Agarwal NJ and Lundquist JK (2026) Characterizing atmospheric stability in complex terrain. *Wind Energy Science* 11: 883-908. DOI: 10.5194/wes-11-883-2026.
17. Becker M, Lejeune M, Chatelain P et al. (2025) A dynamic open-source model to investigate wake dynamics in response to wind farm flow control strategies. *Wind Energy Science* 10: 1055-1075. DOI: 10.5194/wes-10-1055-2025.
18. Doekemeijer BM, Van der Hoek D, Simley E et al. (2021) Field experiment for open-loop yaw-based wake steering at a commercial onshore wind farm in Italy. *Wind Energy Science* 6: 159-176. DOI: 10.5194/wes-6-159-2021.
19. Pandit R, Astolfi D, Hong J et al. (2022) SCADA data for wind turbine data-driven condition/performance monitoring: a review on state-of-art, challenges and future trends. *Wind Engineering* 47(2): 422-441. DOI: 10.1177/0309524X221124031.