

Response to Referee 1

We sincerely thank Referee 1 for the careful and constructive assessment of our manuscript. The referee’s central concern—that the analysis may combine tropical-cyclone wind-field structure with differences between wind products and therefore may not establish the incremental contribution of WWDI* beyond structural contrast and sampling geometry—led us to substantially re-examine the interpretation of the study.

In response, we performed a set of additional analyses specifically designed to address these concerns. These included: (i) a signed decomposition separating the local CYGNSS–ERA5 difference from the ERA5-local-wind versus IBTrACS-storm-maximum difference; (ii) expanded controls for storm structure using normalized radius/RMW and storm-relative asymmetry; (iii) a fixed-algorithm nested feature analysis; (iv) a sensitivity analysis removing the ERA5-wind-dependent wave-age contribution; (v) a sampling-geometry sensitivity analysis for the wave-height-transience component; and (vi) an evaluation of the storm-maximum-targeted adjustment against the local ERA5 wind.

These analyses materially changed the interpretation of the manuscript. The revised manuscript no longer presents WWDI* as an independent local retrieval-error index. Instead, the large CYGNSS–IBTrACS differences are interpreted primarily as reference-scale and storm-structure mismatch, with wind–wave descriptors providing secondary and partly wind-coupled information.

Reviewer Comment 1 (Line 8)

Because the index weights are learned against E_{TC} , WWDI may preferentially select wave variables that proxy the contrast between storm-maximum and local wind. The resulting weights do not, by themselves, establish a contribution specific to CYGNSS retrieval uncertainty.*

Response:

We agree with this concern.

The original WWDI* weights were learned against the absolute CYGNSS–IBTrACS mismatch. Because that target contains the contrast between a local CYGNSS observation and the IBTrACS storm-maximum wind, a component can receive a large weight because it covaries with storm structure or the local-to-maximum contrast, without necessarily containing information specific to local CYGNSS retrieval disagreement.

To test this directly, we performed a fixed-algorithm nested feature analysis in which storm structure, local atmospheric wind, bulk wave variables, wave-state descriptors, and WWDI* were introduced sequentially while the modeling procedure was held fixed.

The results show that local ERA5 wind provides the dominant incremental information after storm intensity and radius are included. Bulk wave variables provide smaller and model-dependent contributions, whereas WWDI* provides only small and inconsistent additional information.

Accordingly, we have revised the manuscript so that the learned WWDI* weights are no longer interpreted as evidence of a contribution specific to CYGNSS retrieval uncertainty. WWDI* is now treated as a secondary descriptor of covarying storm and wind–wave regimes.

Reviewer Comment 2 (Line 12)

The rationale for diagnosing sea-state effects from differences between local CYGNSS wind and storm-maximum IBTrACS wind remains unclear. A substantial difference can arise from cyclone structure even when the local retrieval is accurate.

Response:

We agree. This has become one of the central interpretive changes in the revised manuscript.

We explicitly decomposed the signed CYGNSS–IBTrACS difference into a local CYGNSS–ERA5 component and an ERA5-local-wind versus IBTrACS-storm-maximum component.

Across the four transfer experiments, the RMS magnitude of the local CYGNSS–ERA5 difference is approximately $3.0\text{--}3.2\text{ ms}^{-1}$, whereas the RMS magnitude of the ERA5-local-wind versus storm-maximum term is approximately $13.7\text{--}19.1\text{ ms}^{-1}$. The total CYGNSS–IBTrACS RMS difference is approximately $12.9\text{--}18.1\text{ ms}^{-1}$.

Thus, the large CYGNSS–IBTrACS difference is governed primarily by the contrast between a local wind quantity and a storm-scale maximum rather than by local CYGNSS–ERA5 disagreement.

We have therefore revised the terminology throughout the manuscript. The CYGNSS–IBTrACS quantity is now described as a storm-maximum/reference-scale mismatch, not as a local retrieval error. The CYGNSS–ERA5 quantity is presented separately as a local product-to-product wind difference.

Reviewer Comment 3 (Line 16)

The existing intensity–radius controls are useful, but they do not fully characterize cyclone wind-field structure. Differences in the radius of maximum wind, radial gradients, and quadrant asymmetry could still confound the sea-state relationship, particularly given the reported reversal in correlation sign.

Response:

We agree that storm intensity and simple distance-to-center controls do not adequately describe a strongly structured tropical-cyclone wind field.

We therefore repeated the analysis using expanded structural controls that include normalized radius/RMW and storm-relative asymmetry in addition to storm intensity and radial position.

On the common subset for which the required structural information is available, the partial Spearman association between WWDI* and the absolute local CYGNSS–ERA5 difference is approximately -0.011 , $+0.013$, $+0.009$, and $+0.010$ for WP→BOB, BOB→WP, WP→ATL, and BOB→ATL, respectively.

In contrast, the corresponding associations with the absolute ERA5–storm-maximum structural term remain strong, approximately -0.744 , -0.721 , -0.678 , and -0.663 .

These results strongly support the referee’s concern that the original relationship was primarily organized by storm structure and reference-scale contrast.

We also agree that the reversal of correlation sign should not be interpreted causally. In the revised manuscript, we describe it as evidence of suppression/confounding among storm intensity, radial structure, asymmetry, and wind–wave variables.

A true local radial wind-gradient term cannot be independently reconstructed from the present

one-point matchup dataset. We have therefore stated this explicitly as a remaining limitation rather than implying that the expanded controls fully resolve all aspects of tropical-cyclone wind-field structure.

Reviewer Comment 4 (Line 23)

Reducing mismatch relative to a storm maximum can move peripheral local winds toward that maximum without improving their local accuracy. Improvement over a global-bias baseline therefore appears insufficient to establish a benefit attributable specifically to sea-state information.

Response:

We agree and tested this point directly.

The original adjustment substantially reduces the CYGNSS–IBTrACS storm-maximum mismatch. However, when the adjusted CYGNSS wind is evaluated against the local ERA5 wind, the result is the opposite.

The local RMSE increases from approximately $3.0\text{--}3.2\text{ m s}^{-1}$ before adjustment to approximately $12.6\text{--}15.1\text{ m s}^{-1}$ after the WWDI*-based adjustment across the four transfer experiments.

Thus, the reduction in storm-maximum mismatch does not represent an improvement in local wind agreement. It instead moves local CYGNSS values toward the storm-maximum reference.

We have therefore removed the interpretation of this procedure as an improvement in CYGNSS retrieval accuracy or as evidence of a sea-state-specific correction benefit. In the revised manuscript, it is described as a reference-relative mapping toward storm-maximum intensity.

Reviewer Comment 5 (Line 172)

The manuscript acknowledges the difference in spatial support, but the contributions of the IBTrACS maximum–ERA5 local difference and the CYGNSS–ERA5 local difference are not quantitatively distinguished. Their relative importance is central to interpreting WWDI.*

Response:

We agree and have now quantified this distinction explicitly.

The RMS magnitude of the local CYGNSS–ERA5 difference is approximately $3.0\text{--}3.2\text{ m s}^{-1}$ across the transfer experiments, whereas the ERA5-local-wind versus IBTrACS-storm-maximum term has an RMS magnitude of approximately $13.7\text{--}19.1\text{ m s}^{-1}$.

This magnitude contrast demonstrates that the structural/reference-scale component is substantially larger than the local CYGNSS–ERA5 difference.

The revised manuscript now makes this decomposition a central part of the Results and Discussion rather than treating spatial-support mismatch primarily as a qualitative limitation.

This result also changes the interpretation of WWDI*: its association with the original CYGNSS–IBTrACS quantity is now interpreted primarily in the context of storm structure and reference scale, rather than as evidence of local CYGNSS retrieval uncertainty.

Reviewer Comment 6 (Line 301)

The absolute response obscures whether the storm-maximum-local and CYGNSS-local differences reinforce or offset one another. These two absolute differences are not additive, so associations with E_{TC} alone do not establish their respective contributions.

Response:

We agree.

The original emphasis on absolute mismatch obscured whether the local and structural signed contributions reinforce or oppose one another.

We therefore performed the decomposition first using signed differences and verified the exact algebraic relationship among the total CYGNSS–IBTrACS difference, the local CYGNSS–ERA5 difference, and the ERA5-local-wind versus IBTrACS-storm-maximum difference.

Absolute differences are now analyzed separately only as measures of disagreement magnitude.

The revised manuscript explicitly states that the absolute components are not additive. We therefore do not construct an “absolute-error budget” by summing them, nor do we interpret their RMS values as additive contributions.

This signed-first treatment allows reinforcement and cancellation to be distinguished while preserving the absolute quantities for descriptive mismatch analysis.

Reviewer Comment 7 (Line 306)

E_{ERA5} represents a difference between two wind products rather than an independently observed retrieval error. Even an association with this local discrepancy would not, by itself, attribute the mismatch to CYGNSS or to wave effects.

Response:

We agree completely.

ERA5 is a reanalysis product and is not treated as ground truth in the revised interpretation. Accordingly, the CYGNSS–ERA5 quantity is now described as a local product-to-product wind difference, not as an independently observed CYGNSS retrieval error.

An association with this local discrepancy indicates systematic covariation but does not, by itself, determine whether the difference arises from CYGNSS, ERA5, collocation/representation effects, or a physical sea-state mechanism.

We have revised the manuscript accordingly and removed causal attribution based solely on CYGNSS–ERA5 differences.

We have also distinguished the role of ERA5 atmospheric wind from the separate CMEMS validation of ERA5 significant wave height. The independent CMEMS comparison provides evidence regarding ERA5 SWH representation only; it does not make ERA5 wind an independent truth reference.

Reviewer Comment 8 (Line 319)

ERA5 wind enters the wave-age component as well as the CYGNSS–ERA5 discrepancy. This shared input introduces potential mathematical coupling, which complicates interpretation of any apparent additional information in the index.

Response:

We agree that this shared input can introduce mathematical/statistical coupling.

We therefore performed a sensitivity analysis in which the wave-age component D_A , which contains ERA5 wind, was removed from WWDI*. The originally learned nonnegative weights of the remaining components were renormalized without refitting to the external data.

The resulting no- D_A index remains extremely similar to the original WWDI*, with Spearman correlations of approximately 0.993–0.996 across the four transfer experiments.

After the expanded storm-structure controls, the partial Spearman association between the no- D_A index and the absolute local CYGNSS–ERA5 difference remains weak, approximately +0.017, +0.039, +0.030, and +0.030 for WP→BOB, BOB→WP, WP→ATL, and BOB→ATL, respectively.

Thus, shared ERA5-wind dependence contributes to the coupling and must be acknowledged, but removing the wave-age contribution does not reveal a strong independent local sea-state/mismatch relationship.

We have added this sensitivity analysis to the revised manuscript and explicitly discussed the shared-input limitation.

Reviewer Comment 9 (Line 356)

Differences between successive storm-time medians can reflect changes in sampled radius or quadrant, as well as temporal sea-state evolution. The interpretation of the wave-height-change component as a transient response therefore remains uncertain.

Response:

We agree.

The original SWH-transience descriptor was intended to represent storm-time variation in wave height, but the sampled locations also change as CYGNSS observes different parts of the storm.

We therefore tested the transience descriptor after accounting for sampling geometry using radial standardization and combined radial/quadrant standardization.

The raw agreement of the original T_H descriptor with the spatially standardized versions changes substantially, confirming that the descriptor is influenced by sampling geometry. However, the resulting changes in the overall WWDI* associations are small, with a maximum absolute change in Spearman correlation of approximately 0.02.

We have therefore revised the interpretation of T_H . It is now described as a storm-time SWH transience descriptor rather than as a direct measure of the physical temporal response or lag between wind forcing and wave adjustment.

We have also removed the interpretation that post-peak behavior can be attributed uniquely to delayed wave adjustment, because the present analysis does not independently isolate that mechanism.

Reviewer Comment 10 (Line 648)

Selecting an algorithm separately for each feature group means that performance differences can reflect both the variables and the fitted model. The reported benchmark therefore does not fully isolate the contribution of wave information beyond local wind and cyclone context.

Response:

We agree.

In the original benchmark, separately selecting the best-performing algorithm for each feature group could confound the effect of predictor information with the effect of changing the model class.

We therefore repeated the analysis using fixed algorithms and identical preprocessing across all feature groups.

The fixed-model experiment uses the same modeling procedure while adding information sequentially: storm intensity and radius provide the structural baseline; local ERA5 wind is then added; bulk wave variables and wave-state descriptors are added subsequently; and WWDI* is evaluated only after those predictors are already available.

The resulting hierarchy is consistent across the principal tests. Local atmospheric wind provides the dominant incremental information. Bulk wave variables provide smaller and model-dependent improvements, while WWDI* provides only small and inconsistent additional information.

We have replaced the original algorithm-dependent benchmark interpretation with this fixed-model analysis. The revised manuscript no longer claims that the previous benchmark independently demonstrates a contribution from wave information beyond local wind and cyclone context.

Overall Response and Revised Interpretation

Taken together, the referee's comments and the additional analyses led us to substantially revise the scientific interpretation of the study.

The revised manuscript now states that:

The large CYGNSS–IBTrACS differences observed in tropical cyclones are governed primarily by reference-scale and storm-structure mismatch rather than by local CYGNSS retrieval disagreement. Wind–wave descriptors covary with these storm regimes and provide modest secondary information, but the present evidence does not support WWDI* as an independent local retrieval-error index.

The revised manuscript also makes clear that:

An adjustment calibrated against storm-maximum intensity can substantially reduce the CYGNSS–IBTrACS mismatch while simultaneously degrading agreement with a collocated local-wind reference. Reference-relative mismatch reduction therefore must not be interpreted as retrieval-accuracy improvement.

We sincerely thank the referee for identifying these issues. The additional analyses prompted by the review have led to a clearer distinction among local product disagreement, storm structure, reference-scale mismatch, sampling geometry, and secondary wind–wave-state variability.