Response to comments from reviewer 1

Manuscript: Coastal upwelling and tropical warm water intrusions are key drivers of interannual fog variability along the southwestern African coast # EGUSPHERE-2026-1208

First of all, we would like to thank the reviewer for the positive comments and for providing very helpful and constructive feedback. We thank the four anonymous reviewers in the acknowledgments of the updated version of the manuscript. The comments and suggestions from reviewer 1 are incorporated in italics and addressed by the authors in blue.

1 Major Comment

The CCF input features seem to be an order of magnitude greater than the number of monthly observations for the time-series (thousands of features versus hundreds of observations). The tuning of the lambda was done to maximize out of sample performance, and therefore reduce overfitting. While that can provide an acceptable model for the historical reconstruction analysis in 3.2, it can still lead to many spurious coefficients when working with such a high dimensional problem. Perhaps using a harsher penalty for the sensitivity analysis in 3.1, could yield a more reasonable number of important input features. Alternatively, the authors could also use a dimensionality reduction technique like PCA (i.e. apply PCA to an n x n matrix of spatial points to reduce the total number of input features while conserving spatial variance information). Another option could be to remove the input fields over land for RH at 700 hPa and U10 and V10. This could additionally reduce noise in ridge coefficients because the authors mention that the wind fields are not as reliable over the land as over the ocean. I don't feel super strongly about which approach is used, but I do think the CCF analysis would benefit from some dimensionality reduction or a harsher penalty to clarify which spatial fields and variables control FLC variability.

We agree that a high number of predictors relative to the number of observations can be a concern. The choice of region was made with this in mind, as described in Section 2.3 to balance the size of the region and, therefore, the number of predictors within it, while still capturing all the relevant synoptic-scale controls. Following the reviewer's recommendation, we masked the input fields over land for R700, U10, and V10. We chose this approach because, in addition to helping with dimensionality, it also removes the inland wind fields, which have been shown to not accurately capture the diurnal cycle of the thermo-topographic winds.
It is worth noting that this masking results in a slight decrease in performance for CN-ERA5 (R^2 from 0.32 to 0.30), AN-ERA5 (R^2 from 0.42 to 0.38), and CN-MERRA-2 (R^2 from 0.33 to 0.31), and a slight increase for AN-MERRA2- (R^2 from 0.35 to 0.36). This slight decrease in performance for most configurations likely reflects a trade-off between reducing noise and removing informative signal. Over land, both reanalyses show substantial discrepancies in the wind fields due to complex terrain and differences in model parameterizations, and neither dataset adequately captures the region's thermo-topographic wind circulation. While these inland wind fields are therefore not reliable in an absolute sense, they may still carry some spatially correlated signal related to onshore/offshore flow that the ridge regression exploits, even if imperfectly. Masking these fields removes this signal along with the noise. Overall, we consider this modest trade-off acceptable, since it improves the physical interpretability and robustness of the remaining predictors.
Additionally, Figures 3 and 4 of the manuscript are the most impacted by this change (see comparison in the Appendix of this document).. For Figure 3, the ERA5 and MERRA-2 trends move closer together in the AN, whereas in the CN they diverge slightly; however, in both regions the trends remain statistically not significant. For Figure 4, the wind trend contributions between the two reanalyses also align more closely overall; however, CN-ERA5 remains an outlier with the largest discrepancy, marked by a negative area in the ridge coefficients over the northern part of the region that is present in all four configurations but considerably stronger here. Reducing the domain does not seem like a suitable response, since the pattern is not specific to CN-ERA5. This area is also close to the AN region, and the domain needs to remain large enough around both regions to capture the non-local effects.

- Added in section 2.2 :"To reduce the dimensionality of the predictor set and to avoid relying on inland wind fields, which have been shown to be unreliable (see Fig. B2), the R700, U10, and V10 fields were masked over land, so that only oceanic grid points were retained as predictors for these three variables. This also ensures consistency across all predictors, as they now share the same spatial extent."
- All figures, tables, and their analyses, where necessary, have been updated in the manuscript to reflect this change in method.

2 Minor Comments

It is unclear what the spatial range of AN and CN are, the paper shows the central point, but I didn't see any description of the specific area of focus. I recommend indicating the spatial averaging box for FLC in Figure 1. Also, I recommend using an open circle to show the centers, so that the value for that grid cell is visible and that the color of the circle is not the same as the shades used in the color bar.

The regions are briefly described in the introduction, but we agree with the reviewer that it is necessary to show the actual spatial extent of the regions. The two regions are not outlined in Figure 1, whereas in the rest of the maps their centers are now indicated by open circles.

Ln 129-130: After removing trend and monthly seasonality, is there still autocorrelation in the residuals?

We thank the reviewer for raising this. We tested for residual autocorrelation in the deseasonalized, detrended target time series using the Ljung-Box test (n = 186 months per region), which is significant at lag 12 ($p < 0.001$, both regions). The Autocorrelation function (ACF) (Fig. 1 in this document) shows autocorrelation coefficients exceeding the 95% confidence band near lags 10–12 for both regions, with most shorter lags falling within the band. We do not have a definitive explanation for this pattern; it may reflect multi-month persistence in the large-scale climate drivers of FLC variability discussed in Section 3.3, an incomplete removal of inter-annual shifts in fog-season timing by the fixed monthly climatology, or some combination of factors. We note that our cross-validation uses contiguous, non-shuffled folds, which reduces the potential for train-test dependence relative to randomly shuffled folds, though given the residual autocorrelation at lag 12, some dependence within folds cannot be fully excluded."
- Added in section 2.3: "However, the residual autocorrelation remains in the deseasonalized, detrended monthly time series (Ljung-Box test, $p < 0.001$ at lag 12), concentrated near lags 10–12. It's potentaly explained by the multi-month persistence in the large-scale climate drivers discussed in Section 3.3 or to limitations of the fixed monthly climatology used for deseasonalizing. This may reduce the effective degrees of freedom in subsequent significance tests."

What is the standard deviation in out of sample performance across the cross validation?

Out-of-sample R^2 across the 6 folds (mean $\pm$ std) is 0.36 ± 0.09 (MERRA-2) and 0.38 ± 0.12 (ERA5) for AN, and 0.31 ± 0.26 (MERRA-2) and 0.30 ± 0.26 (ERA5) for CN. AN shows stable skill across both reanalyses, while CN shows substantially greater fold-to-fold variability in both reanalyses, driven by two folds (folds 3–4, covering July 2009–August 2014) with near-zero R^2 in both reanalyses. This period coincides with a period of reduced variance in the observed CN FLC anomalies. Using the CN-ERA5 configuration, we examined this further (Fig. 2 in this document) and found that folds 3–

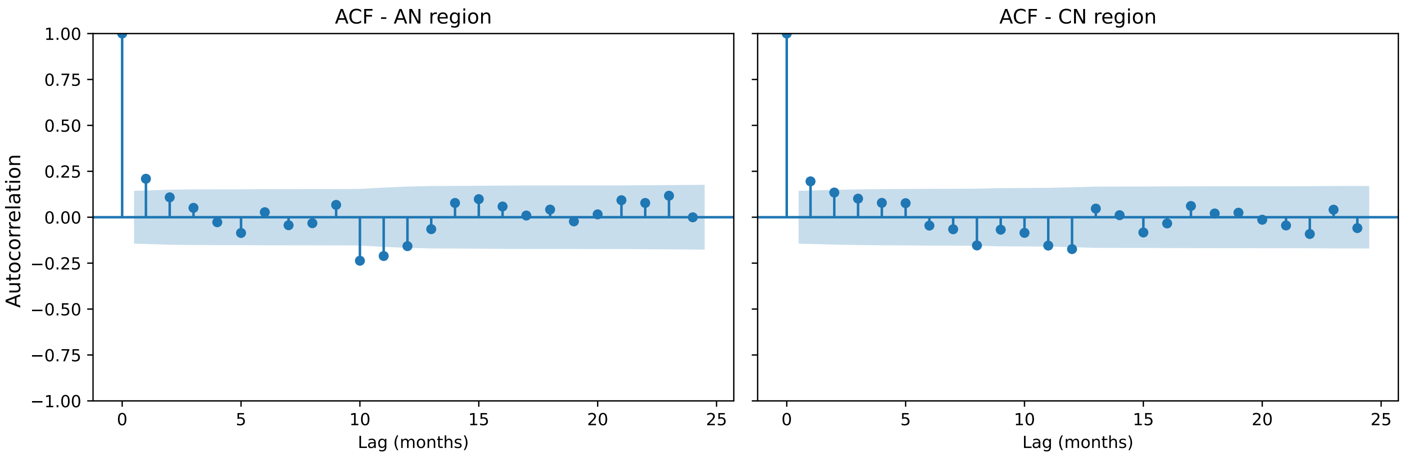

Figure 1: Autocorrelation function (ACF) of the deseasonalized, detrended monthly FLC anomaly time series for the (a) Angolan Namib (AN) and (b) Central Namib (CN) regions (n = 186 months, 2004–2019). Shaded regions indicate the 95% confidence interval under the null hypothesis of no autocorrelation.

4 have markedly lower target variance ($\approx 23\%^2$) than the remaining four folds ($43\%^2$–$58\%^2$). Notably, the model's mean squared error (MSE) in folds 3–4 (23.5$\%^2$–24.3$\%^2$) does not decrease proportionally with this reduced variance, remaining comparable to, or only slightly lower than, MSE values in folds with substantially higher target variance and R^2 (e.g., MSE $= 27.8\%^2$ for a fold with variance $= 43\%^2$ and $R^2 = 0.36$). This suggests that the model exhibits a relatively stable absolute error magnitude across periods with different target variability, resulting in reduced predictive skill for low-variance periods in addition to the mechanical effect of reduced variance on R^2.
- Standard deviation values added in Section 3.1.
- Added in 3.1: "While the AN model shows comparatively stable out-of-sample skill across folds, the CN model exhibits substantially greater fold-to-fold variability in both the ERA5 and MERRA-2 configurations. Two of the six folds, covering July 2009 to August 2014, consistently show near-zero predictive skill regardless of the reanalysis product, coinciding with a period of reduced variance in the observed CN FLC anomalies. The model's absolute error in these folds does not decrease proportionally with the reduced target variability, suggesting a reduced ability to resolve lower-amplitude fluctuations in FLC cover, in addition to the mechanical effect of reduced variance on R^2."

Ln 164: Why only show CN-ERA5? It could be worth discussing one of the AN results as well because some of the findings differ.

We acknowledge that differences exist between the AN and CN sensitivity patterns, largely consistent with the northward shift discussed in Section 3.1 (e.g., the SST sensitivity pattern shifting from the Lüderitz to the Cape Frio upwelling cell). However, these differences do not change the overall picture or conclusions of the study. We therefore prefer to present CN-ERA5 in the main text and keep the remaining configurations (AN-ERA5, CN-MERRA-2, AN-MERRA-2) in the Appendix, as we believe this improves the clarity and focus of the narrative.

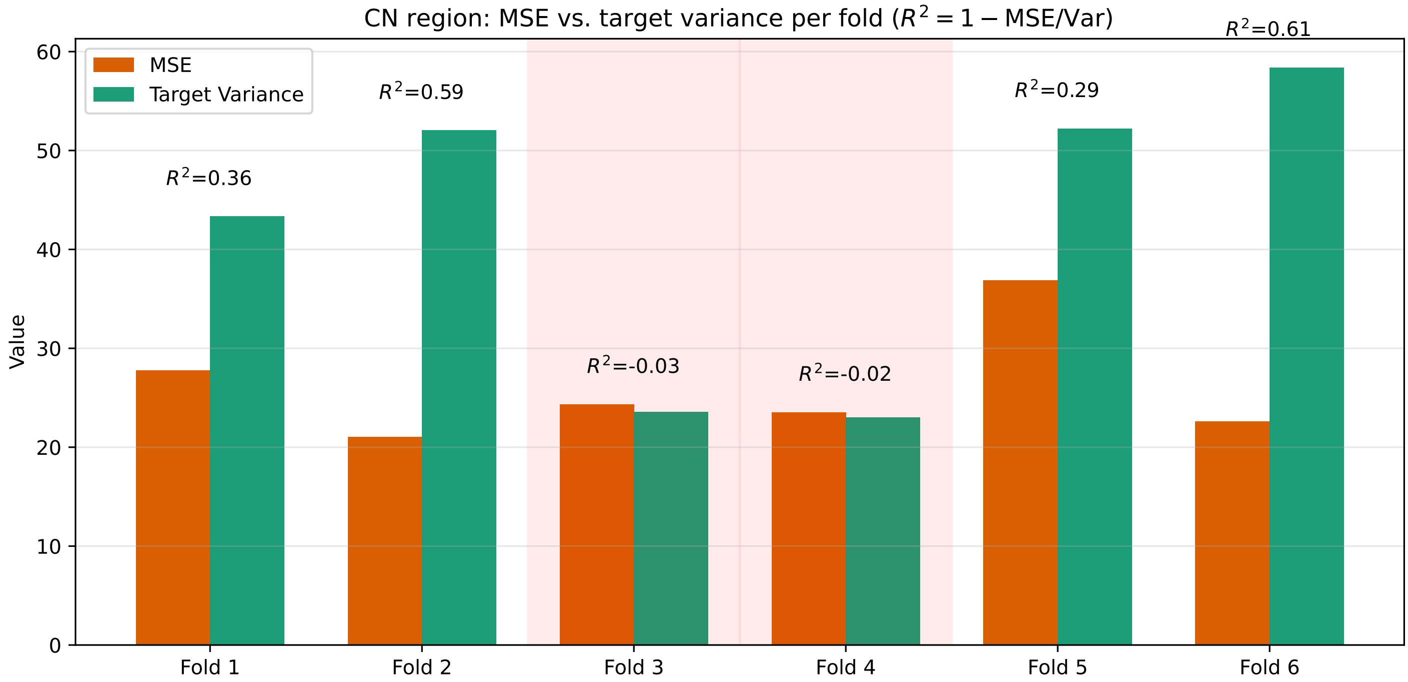

Figure 2: Mean squared error (MSE) and target variance of the observed CN FLC anomaly for each of the 6 cross-validation folds (CN-ERA5 configuration). Folds 3–4 (July 2009–August 2014, highlighted) show target variance roughly half that of the remaining folds, while MSE remains comparable across all folds. R^2 values for each fold are annotated above the corresponding bars.

Ln 190: The negative sensitivity may also be due to the LUCI located to the south of CN.

We considered whether the negative V10 sensitivity could reflect Lüderitz upwelling intensity, given that equatorward (southerly, positive V10) wind is the upwelling-favorable direction along this coast. However, this mechanism would predict a positive, not negative, V10 sensitivity: stronger equatorward wind would intensify upwelling, cool SST south of CN, and increase FLC (consistent with our SST sensitivity), implying a positive association between V10 and FLC through this pathway. Since we observe the opposite sign, we do not believe upwelling intensity via V10 is the primary driver of this sensitivity. This is further supported by the AN configuration, where a similarly negative V10 sensitivity is observed despite the Cape Frio upwelling cell lying much closer to the AN region than the Lüderitz cell lies to CN; if the negative V10 sensitivity were driven by distance-dependent upwelling intensity from a specific upwelling cell, we would expect this sensitivity to differ noticeably between the two regions given their different proximity to their respective upwelling cells. A more consistent explanation is that negative V10 over the region reflects weaker coast-parallel southerly flow, which Andersen et al. (2020) found to be characteristic of fog days in the CN, associated with synoptic-scale disturbances producing a northerly wind anomaly ahead of the trough that favours stratus formation and onshore advection.

Ln 220: How much is the skill improved by averaging annually post-prediction?

We thank the reviewer for prompting us to revisit this. Upon rechecking, we find that annual averaging does not improve skill, contrary to what was stated in the original manuscript (Ln 220); correlation coefficients in fact decrease slightly across all four configurations ($r = 0.66$–0.68 monthly vs. $r = 0.61$–0.64 annual). This discrepancy likely stems from an earlier version of the model used at the time the original text was written. We have therefore modified this sentence in the manuscript.
- Modified in section 3.2: "Averaging the monthly FLC cover anomalies to annual values slightly decreases the correlation with observations ($r = 0.66$–0.68 monthly vs. $r = 0.61$–0.64 annual, across regions and reanalyses)."

Ln 237: What statistical methods are being used to determine the trend? Have the authors considered non-linear trends?

Trends are estimated using ordinary least-squares linear regression (`scipy.stats.linregress`), with statistical significance assessed via the standard OLS p-value. Following this comment, we tested whether a quadratic fit better describes the reconstructed FLC anomaly time series (1982–2019). We compared the two models using the Akaike Information Criterion. For AN, the linear fit was preferred in both reanalyses. For CN, the quadratic fit was modestly favored in both reanalyses. However, a statistically improved quadratic fit does not necessarily correspond to a physically meaningful change in the underlying trend, and we therefore prefer to retain the simpler and consistent linear fit throughout the manuscript.
-Added in Section 2.3: "Trends in the reconstructed FLC anomalies are estimated using ordinary least-squares linear regression. As a robustness check, a quadratic fit to the reconstructed FLC anomaly time series was also tested. The linear fit was preferred for AN, while a quadratic model was slightly favored for CN. However, an improved statistical fit does not necessarily reflect a physically meaningful change in the trend, and therefore the linear fit is retained throughout the manuscript for consistency."

Ln 273: This does make me question why the wind components are in the reconstruction model, maybe justify why it is still important to include them.

Wind components (U10, V10) are retained in the reconstruction model for several reasons. Near-surface wind is a well-established cloud-controlling factor, capturing onshore advection of fog formed over the ocean as well as wind-driven mixing effects on the marine boundary layer (Klein et al., 2018). As noted in our response to the major comment, we have masked land-grid wind values (along with R700), addressing their reduced reliability over land while retaining them over the ocean, where they remain reliable. We find that including these ocean wind components measurably improves the model's predictive skill, supporting their retention as a physically meaningful and useful predictor.
-Added in 2.2: "Finally, U10 and V10 capture near-surface circulation, both the advection of FLCs toward the coast and inland, and boundary layer mixing (Klein et al., 2018)"

Ln 280-281: Another reason the trends are unreliable could be due to the lack of statistical significance as opposed to reanalysis differences.

Added : "The derived trends were found to be unreliable, notably because of discrepancies between ERA5 and MERRA-2, **as well as a general lack of statistical significance in the estimated trends**"

It's a little hard to differentiate the purple and blue on figures 6 and 7, perhaps make the anomalies green as was done in the appendix.

The original color scheme was intended to consistently distinguish ERA5 (purple) from MERRA-2 (green) across figures. We changed ERA5's color to brown throughout the manuscript.

Ln 326:328 I wonder how robust the finding of the correlation between the LUCI/BNI is with the respective AN/CN locations, due to SST anomalies being used to reconstruct the timeseries. I would be curious to see the correlation between any same-sized patch of SST anomalies with the AN/CN anomalies. If the correlation for the LUCI/BNI is significantly higher than any random SST anomaly patch, this could strengthen the findings.

We thank the reviewer for this suggestion. We tested the robustness of the BNI/LUCI–AN/CN correlations by systematically shifting each index's defining box in latitude (northward, southward) and longitude (offshore), and recomputing the correlation with the respective reconstructed AN/CN series at each shift (Figure 3 in this document). For BNI, the correlation remains stable for small shifts and only weakens appreciably once the shift becomes large, consistent with BNI being a spatially broad, literature-defined index (Shannon et al., 1986) situated within one of the strongest SST gradients in the region (Brandt et al., 2024). As BNI is an established index used elsewhere in the literature for comparability, and as even an offshore shift produces only a limited increase in correlation strength, we retain the original BNI definition rather than adjusting its location. For LUCI, the correlation weakens with shifts in most directions, though a slight increase is observed for a northward shift. We attribute this to the fact that upwelling is not confined to the Lüderitz cell: the coastal upwelling band extending north towards the CN region also contributes to FLC variability there, and a northward shift of the box captures it. As this slight increase reflects the contribution of the broader coastal upwelling signal rather than a stronger Lüderitz-specific relationship, and as no shift meaningfully improves on the original definition, we retain LUCI defined directly over the Lüderitz upwelling cell itself (Siddiqui et al., 2023), consistent with its intended physical interpretation
-Added in Section 3.3: "The robustness of the BNI and LUCI definitions to their exact spatial extent was assessed by testing the sensitivity of their correlation with the respective reconstructions to shifts in the defining box location (Appendix C, Figure C3). BNI shows stable correlation for small shifts, weakening only for large shifts, consistent with its broad, literature-defined extent within a strong regional SST gradient. LUCI shows a slight increase in correlation for a northward shift. The Lüderitz cell itself is fixed in location, but upwelling is not confined to it: the coastal upwelling band extending north towards the CN region also matters for FLC variability there, and a northward shift captures it. The increase therefore reflects the broader coastal upwelling signal rather than a stronger Lüderitz-specific relationship."

Also, it would be helpful to mention the periodicity of LUCI/BNI. Does it vary on short or long timescales?

Regarding periodicity, major Benguela Niño events typically persist for six months or longer, with some events (e.g., 1963) lasting close to a full year (Shannon et al., 1986); documented cases include the 1934 event, where the warm water intrusion persisted for about six months, and the 1984 event, associated with a raised sea-level anomaly lasting almost seven months. We note that the Lüderitz upwelling cell relevant to the LUCI index persists year-round as a physical feature, with intensity that is generally weaker in winter and spring (Andrews and Hutchings, 1980; Hutchings et al., 2009); the LUCI index itself, however, represents a deseasonalized SST anomaly and therefore fluctuates around zero rather than reflecting the permanent presence of the underlying upwelling process.
-Added in Secion 3.3: "BNI and LUCI vary primarily on multi-month timescales. Major Benguela Niño events typically

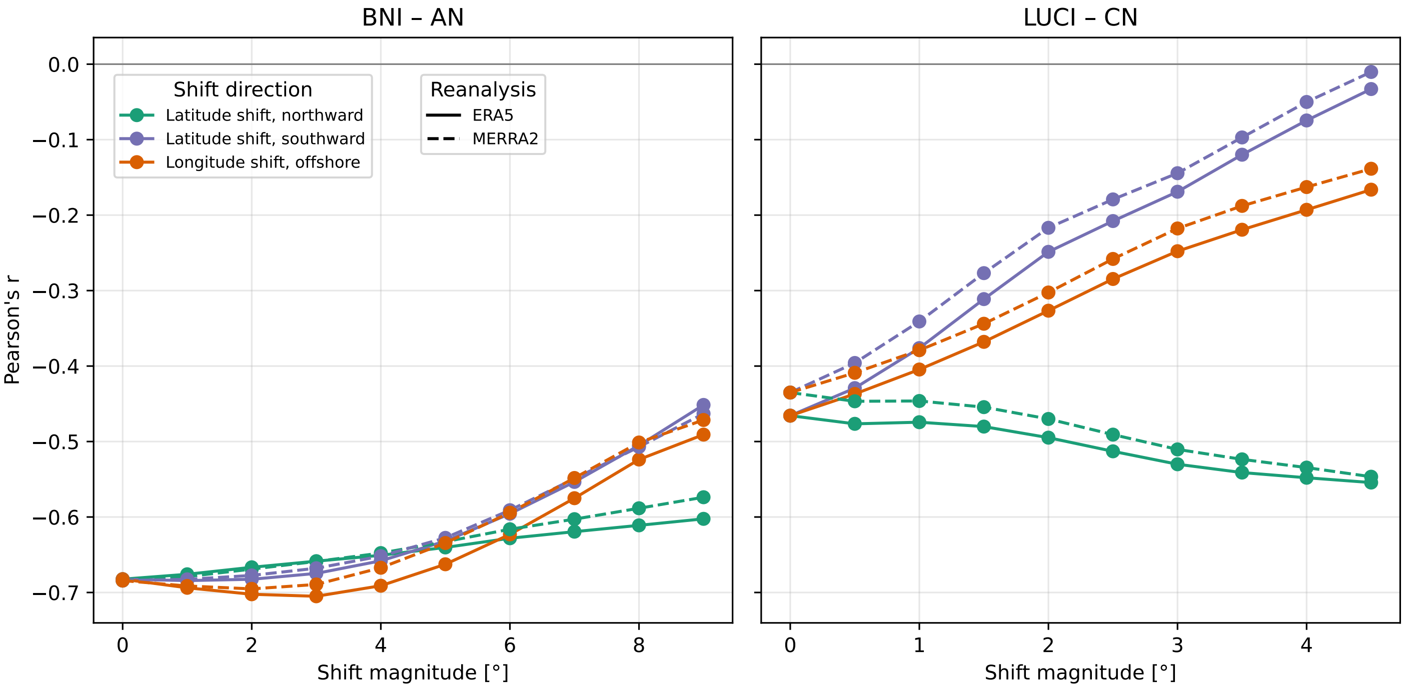

Figure 3: Sensitivity of the Pearson correlation between the BNI/LUCI indices and the reconstructed AN/CN FLC anomalies to shifts in the defining box location. The box is shifted northward, southward, and offshore (westward) from its original location (0° shift), in 1° increments for BNI and 0.5° increments for LUCI, for both the ERA5 (solid) and MERRA-2 (dashed) configurations. **(a)** BNI–AN. **(b)** LUCI–CN.

persist for six months or longer, with some events lasting close to a full year (Shannon et al., 1986). The Lüderitz upwelling cell relevant to the LUCI index persists year-round, with intensity generally weaker in winter and spring (Andrews and Hutchings, 1980; Hutchings et al., 2009). The robustness of the BNI and LUCI definitions to their exact spatial extent was assessed by testing the sensitivity of their correlation with the respective reconstructions to shifts in the defining box location (Fig. C3). BNI shows stable correlation for small shifts, weakening only for large shifts, consistent with its broad, literature-defined extent within a strong regional SST gradient. LUCI shows a slight increase in correlation for a northward shift, likely reflecting partial overlap with the BNI-related tropical warm-water intrusion signal rather than a stronger Lüderitz-specific relationship."
-Added figure C3

Appendix Figure B2

1. *Clouds are most frequent September to March (ln 29), some of these months are missing from the analysis, any reason?*
2. *Does this analysis just compare the closest grid cell to the Marble Koppie station?*

3. Why show the wind direction and not directly the u10 and v10 components?

1. Regarding the missing months: this is a good suggestion; we had not previously examined the diurnal wind direction cycle beyond DJF and JJA. The original focus on DJF and JJA was motivated by the contrasting summer and winter diurnal cycles described for the central Namib by Lindesay and Tyson (1990), which our comparison aimed to test the reanalyses against. We have added a second row to Fig. B2 (panels d–f) covering MAM and SON. The results show a similar bias between the station observations and the reanalyses as seen in DJF/JJA, though with a larger overlap in the interquartile ranges for MAM/SON compared to the clearer separation seen in DJF/JJA.

2. Regarding the grid cell used: yes, this analysis compares the nearest ERA5 and MERRA-2 grid cells to the Marble Koppie station. We have added the exact coordinates of the station and the two nearest grid cells to the figure.

3. Wind direction was originally shown as a simplified, easily interpretable summary of the wind field, without a wind speed dimension. We agree this does not fully convey the u10/v10 information and, in response, have added seasonal wind roses (Fig. B3) for the station and both reanalyses, showing the joint distribution of wind direction and speed. Together with the diurnal wind direction cycle (Fig. B2), this now provides a more complete picture of the wind field, including a strong bias between the station observations and the reanalyses that is consistent across both figures.

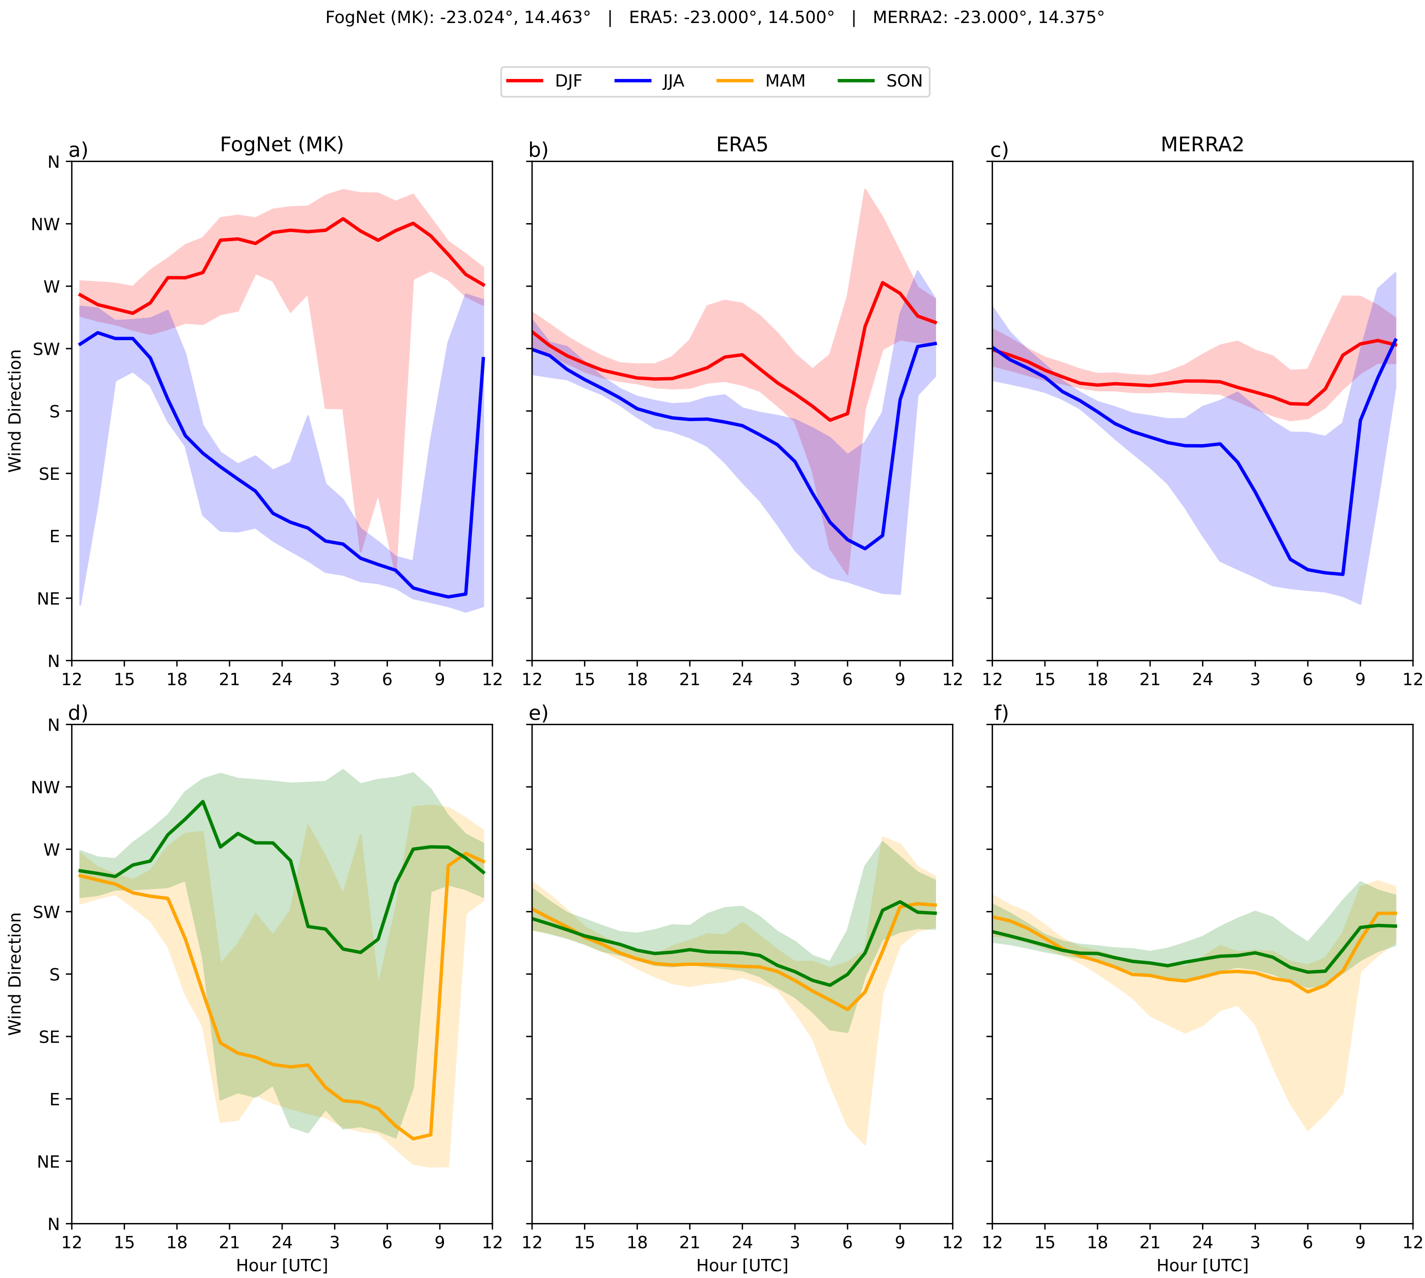

Figure 4: Seasonally averaged diurnal wind direction cycle for 2016 (red: DJF; blue: JJA; orange: MAM; green: SON). Shaded areas indicate the interquartile range (25th–75th percentile) for each hour. Top row: **(a)** FogNet station (Marble Koppie: 23.024°S, 14.463°E), **(b)** ERA5 (nearest grid cell), **(c)** MERRA-2 (nearest grid cell). Bottom row: **(d–f)** as in (a–c), for MAM and SON.

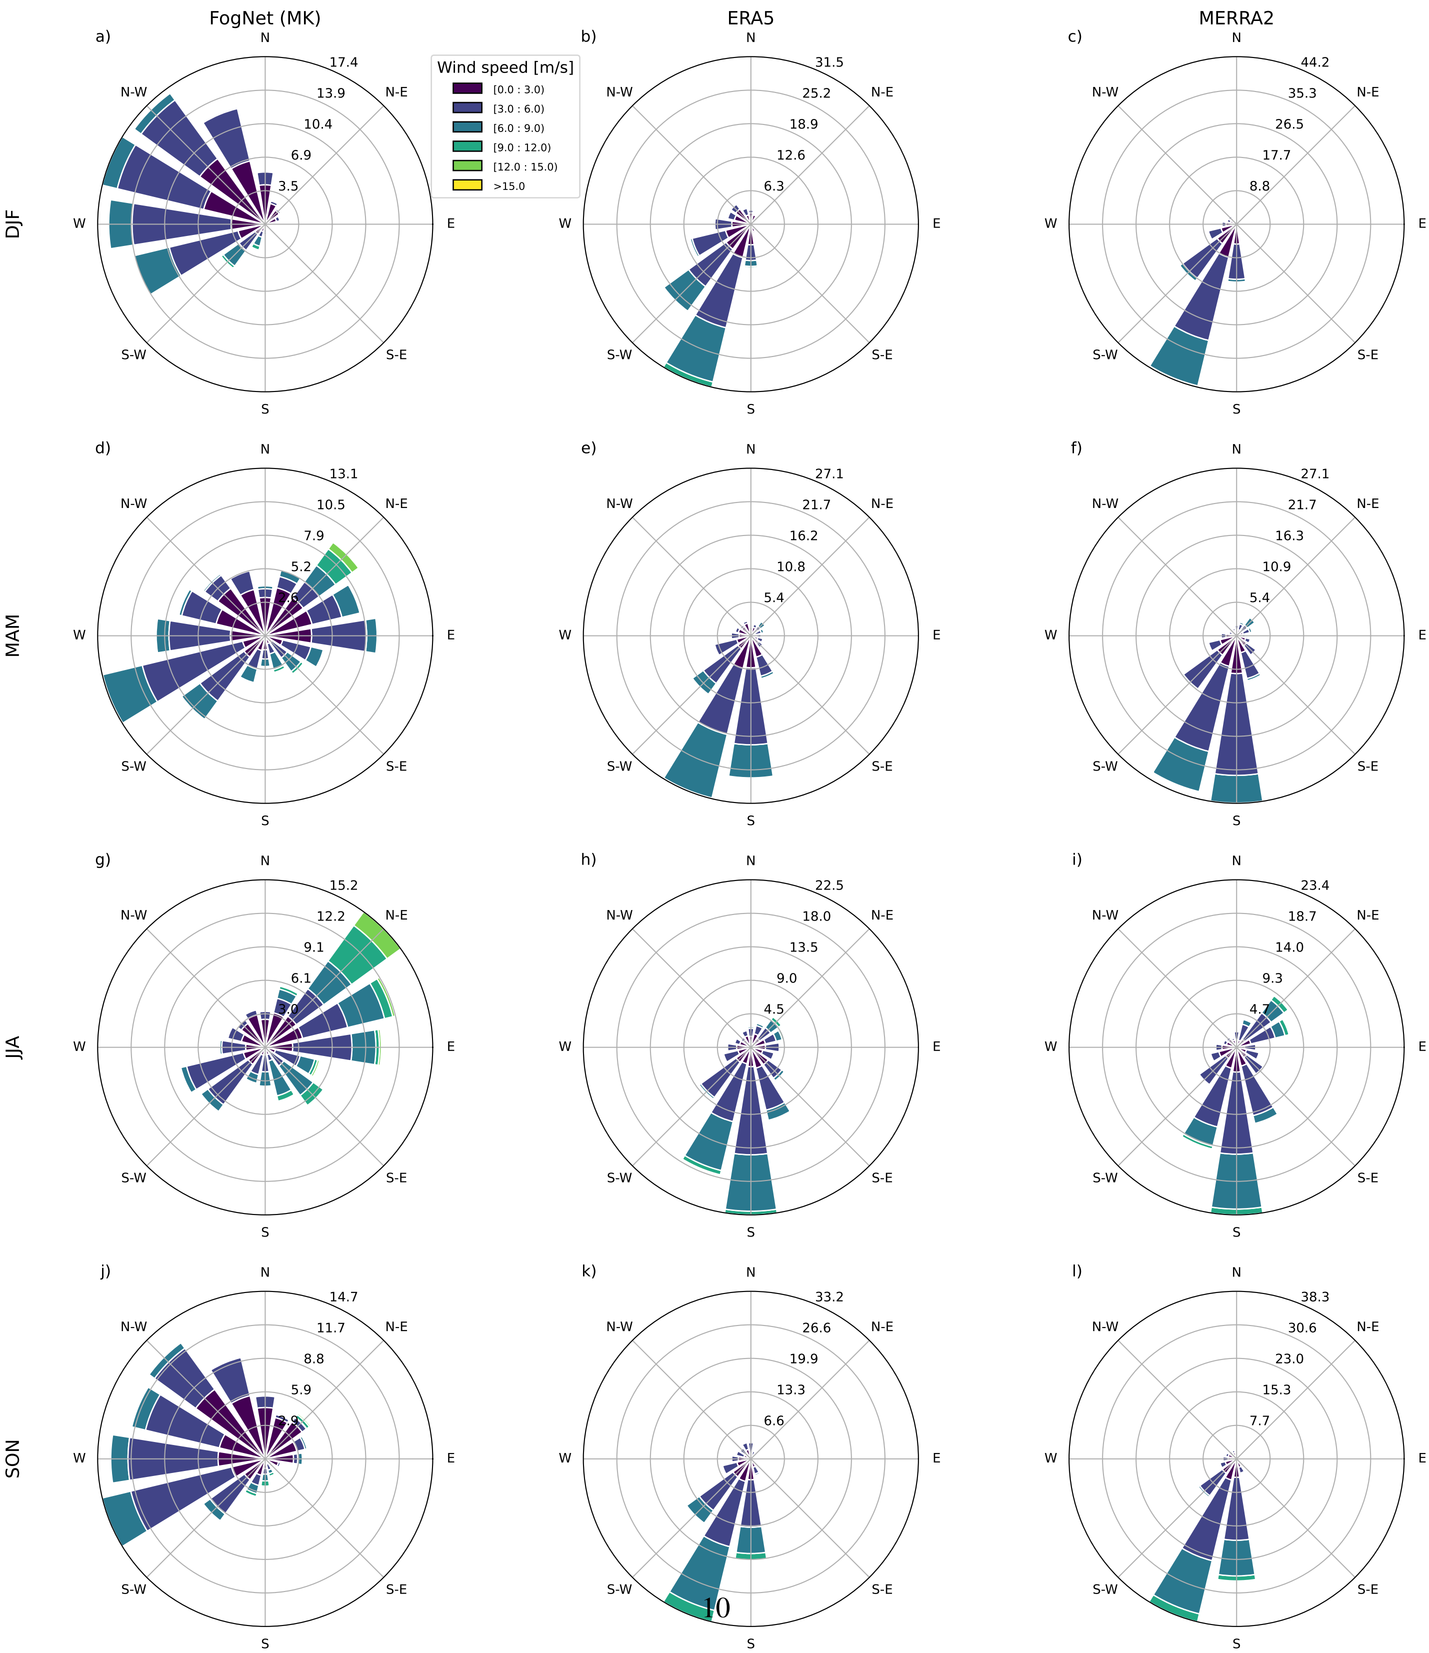

Figure 5: Wind roses showing the joint distribution of wind direction and speed for the FogNet station (Marble Koppie), ERA5, and MERRA-2, separated by season (rows: DJF, MAM, JJA, SON) for 2016. Bars indicate the frequency of occurrence (%) for each wind direction sector, colored by wind speed bin (m/s).

A Old vs New Figures 3 & 4

References

Andersen, H., Cermak, J., Fuchs, J., Knippertz, P., Gaetani, M., Quinting, J., Sippel, S., and Vogt, R.: Synoptic-scale controls of fog and low-cloud variability in the Namib Desert, Atmospheric Chemistry and Physics, 20, 3415–3438, https://doi.org/10.5194/acp-20-3415-2020, 2020.

Andrews, W. and Hutchings, L.: Upwelling in the Southern Benguela Current, Progress in Oceanography, 9, 1–81, https://doi.org/https://doi.org/10.1016/0079-6611(80)90015-4, 1980.

Brandt, P., Bordbar, M. H., Coelho, P., Koungue, R. A. I., Körner, M., Lamont, T., Lübbecke, J. F., Mohrholz, V., Prigent, A., Roch, M., et al.: Physical drivers of Southwest African coastal upwelling and its response to climate variability and change, in: Sustainability of southern African ecosystems under global change: Science for management and policy interventions, pp. 221–257, Springer, https://doi.org/10.1007/978-3-031-10948-5$_9$, 2024.

Hutchings, L., van der Lingen, C., Shannon, L., Crawford, R., Verheye, H., Bartholomae, C., van der Plas, A., Louw, D., Kreiner, A., Ostrowski, M., Fidel, Q., Barlow, R., Lamont, T., Coetzee, J., Shillington, F., Veitch, J., Currie, J., and Monteiro, P.: The Benguela Current: An ecosystem of four components, Progress in Oceanography, 83, 15–32, https://doi.org/https://doi.org/10.1016/j.pocean.2009.07.046, eastern Boundary Upwelling Ecosystems: Integrative and Comparative Approaches, 2009.

Klein, S. A., Hall, A., Norris, J. R., and Pincus, R.: Low-cloud feedbacks from cloud-controlling factors: A review, Surveys in Geophysics, pp. 135–157, https://doi.org/10.1007/s10712-017-9433-3, 2018.

Lindesay, J. A. and Tyson, P. D.: Thermo-topographically induced boundary layer oscillations over the central Namib, southern Africa, International Journal of Climatology, 10, 63–77, https://doi.org/10.1002/joc.3370100108, 1990.

Shannon, L., Boyd, A., Brundrit, G., and Taunton-Clark, J.: On the existence of an EI Niño-type phenomenon in the Benguela System, Journal of Marine Research, 44, 495–520, URL https://elischolar.library.yale.edu/journal$_{o}f_{m}arine_{r}esearch/1826, 1986$.

Siddiqui, C., Rixen, T., Lahajnar, N., Van der Plas, A. K., Louw, D. C., Lamont, T., and Pillay, K.: Regional and global impact of CO2 uptake in the Benguela Upwelling System through preformed nutrients, Nature Communications, 14, 2582, https://doi.org/10.1038/s41467-023-38208-y, 2023.

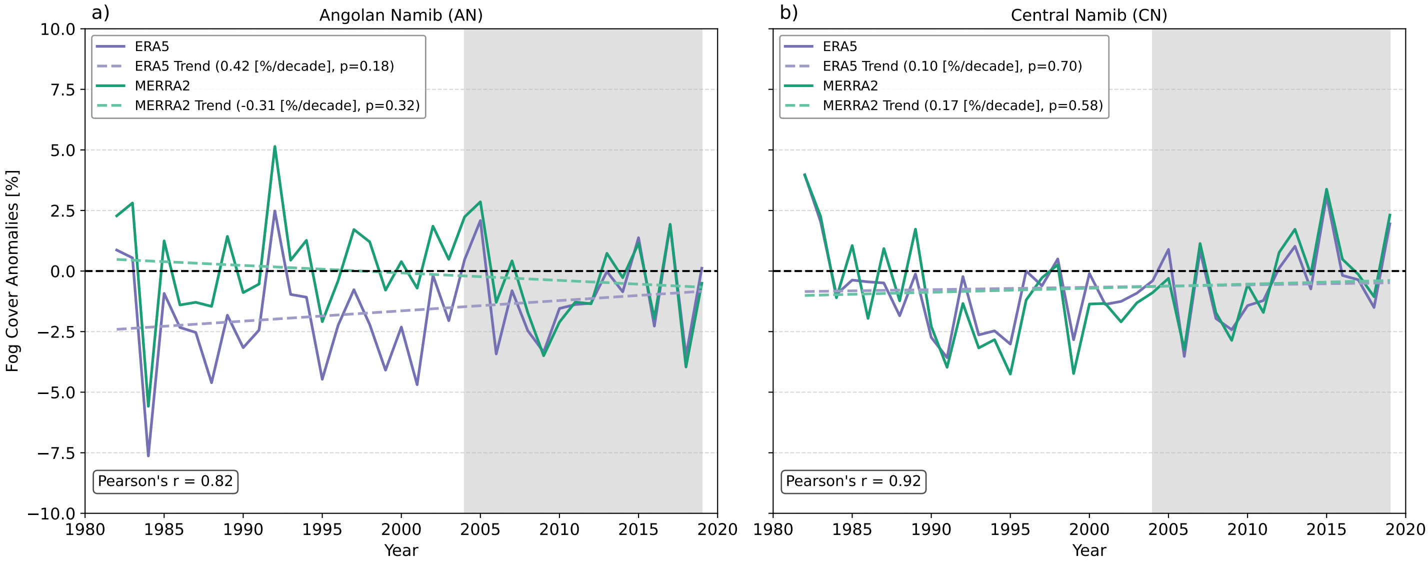

Figure 6: Old version of Figure 3 (with R700, V10 and V10 predictors field not masked over land).

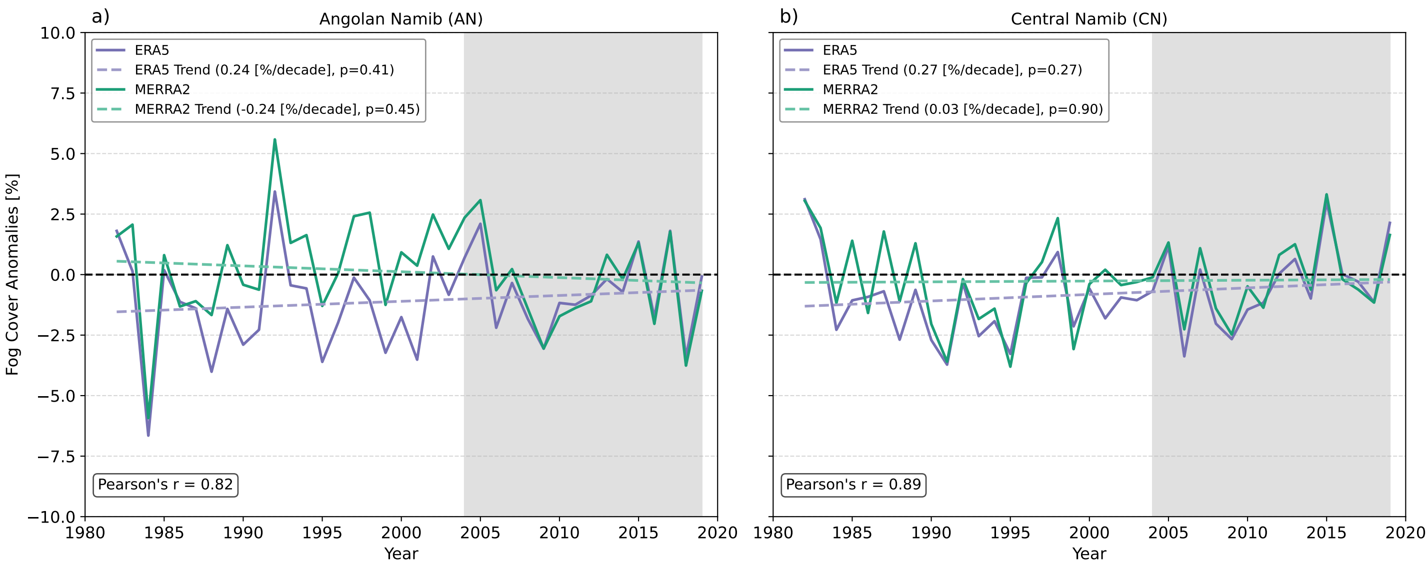

Figure 7: New version of Figure 3 (with R700, V10 and V10 predictors field masked over land).

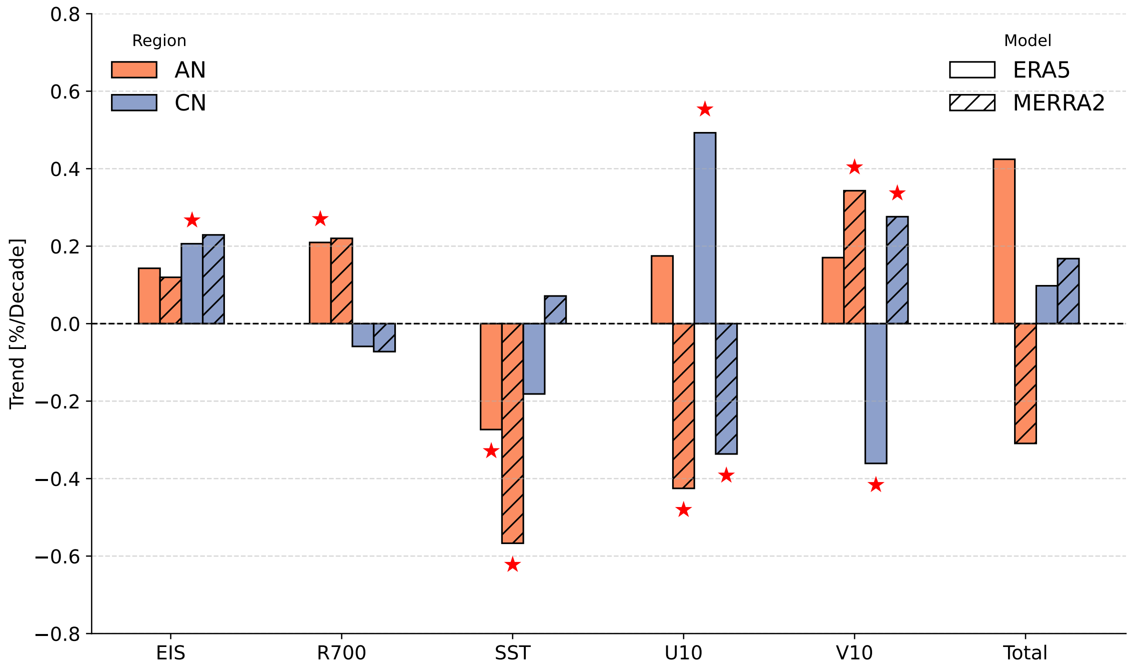

Figure 8: Old version of Figure 4 (with R700, V10 and V10 predictors field not masked over land).

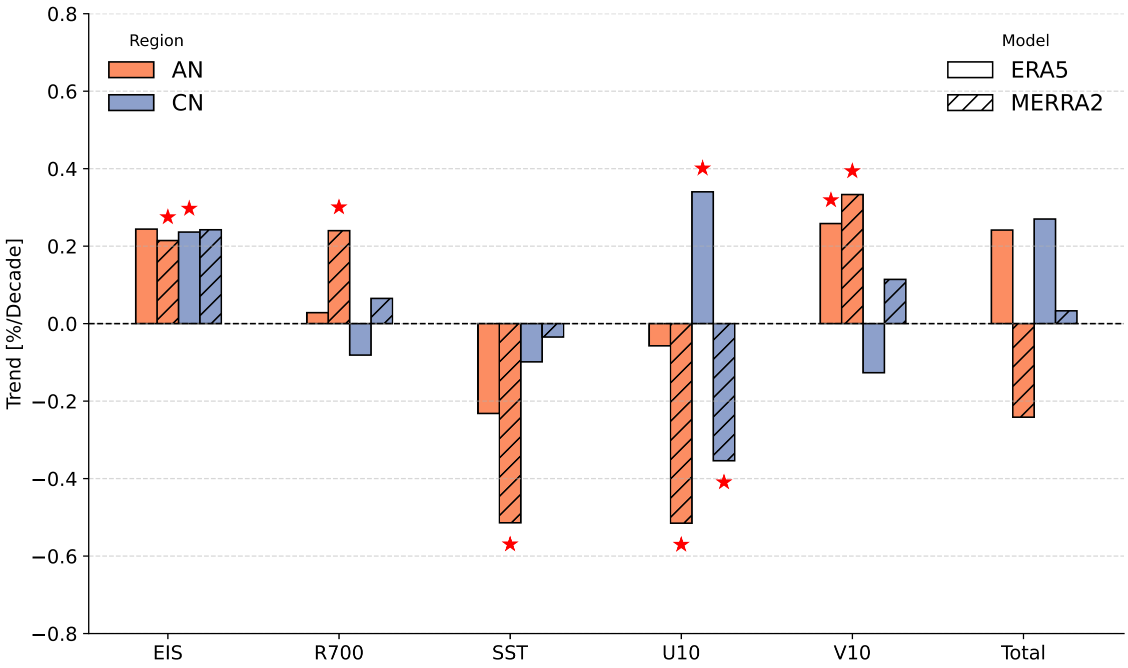

Figure 9: New version of Figure 4 (with R700, V10 and V10 predictors field masked over land).