Review of the manuscript

“Composite Detection of Continental Shelf Fronts from Sea Surface Temperature and Altimetry: Assessing the Added Value of SWOT KaRIn”

This manuscript examines shelf fronts in the Southwestern Atlantic using a composite SST-ADT approach, with particular attention to the contribution of SWOT. I found the comparison across different frontal environments interesting, especially the contrasting results for the shelf-break and coastal fronts. The study addresses a relevant question for Ocean Science, and the workflow is easy to follow. My main concern is how the differences in joint front probability are interpreted. The results show that some products more frequently produce ADT gradients near SST fronts, but I am less convinced that this alone demonstrates better front detection, or that the differences can be attributed specifically to KaRIn. Some additional analysis would help here, together with a more careful description of what the metric measures. Several methodological details also need clarification. I think the paper has merit, but I would recommend major revisions before publication. My main comments are below.

Major comments

1. Sections 2.2.2 and 4; Fig. 12: As I understand Eq. (6), P_joint measures how often the two gradient criteria are satisfied within the chosen spatial tolerance. It does not directly measure detection accuracy. For example, a broader or more persistent area of elevated ADT gradients could increase P_joint without necessarily improving the location of a front. The agreement with previously documented frontal regions is encouraging, although I do not think it is sufficient to infer that false detections are uncommon, as suggested in lines 200-208.

It would help to see the SST-only, ADT-only, and joint occurrence fields together, so that the reader can assess what the composite criterion adds and which features it excludes. For a few representative fronts, a comparison with independent observations, where available, would make the interpretation more convincing. A controlled test with known front locations could also be informative. I am not suggesting that an extensive new validation dataset is essential, but some assessment beyond maximum P_joint seems important if improved detection is to remain a central claim. If independent validation is not feasible, a more limited interpretation in terms of frontal co-occurrence would still be useful.

2. Section 2.1.1; lines 393-402: I was not entirely clear whether CMEMS and SWOT MIOST differ only in the inclusion of KaRIn, or whether there are also relevant differences in their mapping and upstream processing. This matters for the paper's main conclusion. A compact

table giving the exact product versions, contributing missions, DT/NRT status, mapping method, MDT, and corrections would help. It would also be useful to confirm explicitly that the CMEMS release used here contains no KaRIn observations.

Ideally, the comparison would use otherwise matched reconstructions with and without KaRIn. I appreciate that these may not be available to the authors. In that case, I would suggest acknowledging the other processing differences and being more cautious about assigning the full difference to SWOT. Ballarotta et al. (2025) “Integrating wide-swath altimetry data into Level-4 multi-mission maps” seems particularly relevant to this discussion.

3. lines 145-153; Appendices A and B: The percentile approach is straightforward, but I had some difficulty determining exactly which samples enter the threshold calculation. Are the percentiles calculated over the entire space-time domain, and are they then held fixed for the seasonal analyses? Is the same spatial mask used for all products? Giving the actual thresholds in physical units, along with their calculation domain, would make the comparison easier to assess.

I am not fully persuaded by the narrower dynamic range of ADT gradients as the explanation for using a lower percentile, since rescaling gradient magnitudes alone would not change their percentile rank. The choice may well be reasonable in practice, but its justification would be stronger if linked to the detection results. Fig. B1 is helpful in this respect. However, it concerns SWOT MIOST in DJF, whereas the broader conclusion depends on differences among all four products. I would suggest extending the sensitivity analysis to check whether the main product comparisons remain similar over a reasonable range of thresholds and matching radii.

4. Eqs. 3-6; Fig. 12: A little more detail on the transition from the 0.05-degree OSTIA grid to the 0.125-degree analysis grid would be helpful. In particular, are SST gradients computed before or after remapping? How are coastal boundaries and missing values treated? The physical extent of the matching neighbourhood also varies with latitude, which is worth making explicit.

I noticed a difference between Eq. (6), which uses a fixed T, and the public example notebook, where seasonal probabilities are divided by valid-sample counts while the annual plot uses a direct mean of the binary mask. This may not affect the results where coverage is complete, but it would be helpful to clarify which denominator was used in the manuscript, especially for the swath analysis, and how validity is handled when an SST match lies in a neighbouring cell.

5. Section 4.6; Fig. 11: There appears to be a date inconsistency: Section 2.1.2 gives March-July 2024 for Cal/Val, whereas Section 4.6 gives March-July 2023. I assume this is a typographical error, but the actual dates and cycles should be clear. The phrase “filtered SSH + MDT” would also benefit from the exact variable name, since SSH and sea-level anomaly are not equivalent here.

The L3 results are promising, but I found the processing description too brief to judge the comparison fully. My concern is that larger L3 gradients could reflect both finer ocean signals and residual errors, particularly near the coast in this macrotidal region. A representative matched transect, perhaps accompanied by a check of sensitivity to filtering or quality control, would help support the interpretation.

6. Sections 4.2-4.5 and Conclusion: The discussion of regional mechanisms is useful, and I agree with the caution expressed in lines 240-244 about inferring causality. Co-located SST and ADT gradients are consistent with a frontal circulation, but do not uniquely identify a baroclinic density front or its forcing. Salinity compensation and the contribution of background or barotropic flow may also matter. Finally, Fig. 9 marks a transect, but I could not find the corresponding profiles. These would be useful for assessing the peak around 100 km offshore and the claim of improved frontal sharpness (lines 259-265 and 408-410).

Specific comments

1. The abstract, lines 15-17, describes the Cal/Val application as future work, although it is already included in Section 4.6. The wording could be brought into line with the analysis presented.

2. Fig. 4 appears to compare products in DJF, rather than different seasons as stated in the caption. The description of the shelf mask as “continental shelf (outside the 200-m isobath)” is also confusing. The reference to Fig. 4 as an annual mean in lines 228-229 seems to need correction.

3. A latitude value is missing in line 220, after “South of approximately”.

4. The statement that GLORYS12v1 does not capture the San Matías front (lines 248-250) seems too categorical given the reported maximum of 35.16%. Would “a weaker or less complete representation” describe the result more accurately?

5. The third row of Fig. B1 is not explained in the caption. Identifying all three regions and

the product shown would help.

6. For Fig. 10 and the summary statistics, it would be useful to indicate where the quoted maximum occurs and whether it belongs to the Magellan plume front or the Tierra del Fuego tidal front. I was not always sure how the two features were separated in the comparison.

7. Some labels and annotations are difficult to read, particularly in Figs. 11 and B1. Slightly larger text and consistent labelling of probabilities versus percentages would improve readability.

8. I suggest removing the title “Bathymetry of the Study Area” from Figure 2.