Comment 1 – Justification of the adaptive grid design and selection of the number of centroids

Reviewer comment

The manuscript introduces an adaptive grid based on balanced K-means clustering and presents results using configurations with 1300 and 2300 centroids (figure 4, lines 315-322). However, the rationale behind the selection of these values is not sufficiently explained. The authors should explain clearly in the text:

- Why the 1300-centroid configuration was determined.
- Whether any optimization or sensitivity analysis was performed.
- The relationship between the number of centroids, the number of radials per cell, GDOP distribution, spatial coverage, and reconstruction skill.
- The practical criteria that would allow another HF radar operator to select an appropriate number of centroids for a different network configuration.

An additional schematic figure illustrating the balanced K-means procedure and the generation of the adaptive grid would significantly improve readability and reproducibility. This could be added in appendix A section.

Author response

We thank the reviewer for this constructive comment.

The number of centroids represents the main user-defined parameter controlling the adaptive discretisation and should not be interpreted as an optimisation parameter in the classical sense. Rather, it plays a role analogous to the grid spacing in conventional regular-grid products, controlling the trade-off between effective spatial resolution and observational support.

The 1300-centroid configuration was selected to provide a number of adaptive reconstruction nodes comparable to the number of cells adopted in the regular-grid configuration. This choice was made to enable a direct comparison between the conventional fixed-grid representation and the adaptive discretisation while maintaining a similar order of magnitude in the number of reconstruction elements.

No formal optimisation procedure was performed to identify an optimal number of centroids. This was a deliberate choice, as the appropriate discretisation depends on the characteristics of the radar network and on the intended application. As shown by the comparison between the 1300- and 2300-centroid configurations, increasing the number of centroids results in smaller adaptive cells and higher spatial resolution, particularly in densely sampled coastal regions, while reducing the number of radial observations associated with each reconstruction node. Conversely, a smaller number of centroids increases observational support at the expense of spatial detail.

The relationship between the number of centroids, radial support, spatial coverage, and reconstruction geometry is illustrated in Figs. 3 and 4. The constrained clustering redistributes nodes according to the density of radar intersections, maintaining a relatively homogeneous number of observations per reconstruction node across the domain despite the strong offshore decrease in radar sampling density. For the configuration adopted in this study, the average number of radials per cell remains approximately constant at about 4–5 observations per node.

For application to other HF-radar systems, the centroid number should therefore be selected according to the desired balance between spatial resolution and observational support, considering factors such as the density of radial intersections, the expected geometric conditioning of the reconstruction, and the scientific or operational objectives of the final product. In this sense, the parameter should be regarded as a discretisation choice rather than a quantity possessing a universal optimum value.

Furthermore, the implementation developed in this study has been designed to be fully scalable and easily transferable to other HF-radar systems. The grid-generation software, which will be made available to the HF-radar community, allows the number of centroids to be modified straightforwardly by the user. This enables operators to adapt the discretisation to the characteristics of their own network, considering factors such as radar geometry, intersection density, observational coverage, and the desired balance between spatial resolution and observational support. Consequently, the methodology is not tied to the specific configuration analysed in this work, but can be readily applied to different domains and network geometries without modification of the underlying algorithmic framework.
We agree that this aspect is central for the reproducibility and transferability of the methodology. To avoid possible ambiguity, we will make the rationale behind the choice of the centroid number and its practical interpretation more explicit in the revised manuscript.

Comment 2 – Limited quantitative evidence of performance improvement

Reviewer comment
The manuscript convincingly demonstrates that the adaptive grid provides a different spatial representation and resolution of the radar observations. However, the validation results show only limited differences in RMSE and correlation when compared with drifter observations.
The authors should better justify why the proposed method represents a significant advancement despite the relatively small improvements in traditional validation metrics.

Author response
We agree that the adaptive framework does not produce large changes in conventional pointwise validation metrics such as RMSE and correlation. However, this behaviour is consistent with the objective of the proposed methodology and with the nature of HF-radar uncertainty.
The primary aim of the adaptive discretisation is not to improve the accuracy of the underlying radial measurements nor to modify the vector reconstruction algorithm itself. The same radial observations, quality-control procedures, and weighted least-squares reconstruction are used in both the regular-grid and adaptive-grid configurations. Consequently, the dominant sources of reconstruction uncertainty—namely measurement noise, radar geometry, and spatial sampling limitations—remain largely unchanged.
The main contribution of the proposed approach is instead the introduction of an observation-driven discretisation framework that explicitly accounts for the spatial variability of radar sampling geometry. Conventional regular grids impose a uniform spatial discretisation regardless of the distribution of observational information. In contrast, the adaptive framework redistributes reconstruction nodes according to the local density and geometry of radar observations, providing a representation that is more consistent with the intrinsic information content of the observing system.
This distinction is reflected in several results presented in the manuscript. The adaptive grid increases the utilisation of valid reconstruction nodes, improves spatial coverage, reduces radar–drifter separation distances, and provides a more continuous representation of the velocity field, particularly in coastal and geometrically complex regions. The analysis of radar–drifter separation distances further indicates that part of the observed scatter in validation results is associated with spatial representativeness effects rather than with the intrinsic accuracy of the radar measurements themselves.
From this perspective, the proposed framework should be viewed less as an optimisation of the reconstruction algorithm and more as a change in how HF-radar observations are represented spatially. The scientific relevance of the method lies in the transition from a fixed discretisation to an information-adaptive discretisation, in which the spatial resolution reflects the local observational support available in the radar measurements.

This characteristic is particularly relevant for applications that depend strongly on the spatial structure of the reconstructed flow field, such as the identification of transport pathways, mesoscale and submesoscale circulation features, long-term climatological analyses, data assimilation, and operational coastal monitoring. In these contexts, the spatial representativeness and consistency of the reconstructed field are often as important as traditional pointwise validation statistics.

Comment 3 – Spatial representation of validation results

Reviewer comment

Figure 7 presents the dependence of RMSE on GDOP classes. While useful, it does not show where these differences occur spatially.

I recommend adding a spatial map of reconstruction errors (or differences between adaptive and regular grids) to identify:

- Regions where the adaptive grid performs better.
- Coastal versus offshore behaviour.
- The relationship between error patterns and radar geometry.
- Areas where the adaptive approach provides the largest benefit.

A spatially explicit analysis would considerably strengthen the interpretation of the results.

Author response

Spatially explicit analyses can certainly provide additional insight into the geographical distribution of the reconstruction characteristics. In the present study, however, the focus was not on identifying localised improvements in pointwise error metrics, but rather on assessing how different discretisation strategies influence the spatial coverage, representativeness, and organisation of the reconstructed current fields.

A direct spatial comparison of radar–drifter reconstruction errors is challenging because the radar–drifter collocations are irregularly distributed in space and time and are strongly influenced by the drifter trajectories themselves. Consequently, spatial maps of reconstruction error tend to reflect the sampling characteristics of the validation dataset as much as the performance of the reconstruction methodology. For this reason, the manuscript places greater emphasis on analyses that directly address the consequences of the adaptive discretisation. These include the comparison of effective node utilisation (Fig. 6), the dependence of reconstruction error on radar geometry through GDOP (Fig. 7), the spatial continuity of the reconstructed velocity fields (Figs. 8 and 9), and the radar–drifter separation analysis (Fig. 10). Together, these results provide a spatially explicit assessment of the areas where the adaptive framework modifies the representation of the observations and where its benefits are most evident, namely regions characterised by heterogeneous sampling density, coastal constraints, and complex radar geometry.

We also note that the principal advantage of the adaptive framework is not a local reduction in pointwise reconstruction error, but a more observation-consistent representation of the flow field. Accordingly, the areas where the adaptive approach provides the largest benefit are primarily identified through improved spatial coverage, increased effective node utilisation, and reduced representativeness mismatch rather than through large changes in conventional RMSE values.

To clarify this point, we will strengthen the discussion linking the spatial analyses presented in Figs. 6–10 with the interpretation of the validation results, making the geographical context of the observed improvements more explicit in the revised manuscript.

Comment 4 – Spatial representativeness analysis

Reviewer comment

The discussion associated with Figure 10 is one of the strongest parts of the manuscript because it directly addresses representativeness errors in radar–drifter comparisons.
To reinforce this result, I suggest adding a spatial map showing the distribution of radar–drifter separation distances (or the reduction achieved by the adaptive grid relative to the regular grid). Such a figure would help the reader understand where the representativeness gains occur and how they relate to HF radar coverage geometry.
This analysis would also provide additional support for the authors' argument that the adaptive grid improves the spatial consistency of radar-derived current products.
The findings shown in Figure 10 also deserve a stronger discussion. The reduced radar–drifter separation is potentially important for several downstream applications, including Lagrangian transport analyses, search and rescue modelling, oil-spill forecasting, data assimilation and validation of coastal circulation models. The practical implications of this improvement should be emphasized more clearly.

Author response

The radar–drifter separation analysis provides a direct measure of representativeness effects in radar–drifter comparisons. One of the main motivations behind the proposed framework is to replace a fixed spatial discretisation with a representation that adapts to the local distribution of observational information. The reduction in radar–drifter separation distances observed for the Voronoi-based grid is consistent with this objective, indicating a closer correspondence between the reconstruction nodes and the geometry of the observations.
Regarding the proposed spatial representation of radar–drifter separation distances, we acknowledge that such a figure could provide additional geographical context. However, as with reconstruction errors, the spatial distribution of collocations is largely controlled by the drifter trajectories and is therefore highly irregular in both space and time. For this reason, we focused on statistics that quantify the overall representativeness mismatch and on spatial analyses that illustrate where the adaptive framework modifies the effective distribution of reconstruction nodes and spatial coverage.
In particular, the combination of Figs. 6, 8, 9, and 11 already provides spatially explicit evidence of where the adaptive discretisation has the largest impact. The strongest effects are observed in coastal regions and in areas characterised by heterogeneous radar sampling density and complex intersection geometry, where the redistribution of reconstruction nodes leads to improved spatial coverage and reduced representativeness mismatch. The enhanced radar–drifter correspondence observed in Fig. 11 is consistent with this interpretation.
We fully agree with the reviewer that the practical implications of the reduced radar–drifter separation deserve greater emphasis. Beyond validation exercises, improved spatial representativeness is expected to benefit applications that depend strongly on the spatial coherence of the reconstructed flow field, including transport pathway identification, Lagrangian analyses, coastal circulation studies, data assimilation, and operational ocean-monitoring activities. We will strengthen the discussion of these aspects in the revised manuscript in order to better highlight the broader implications of the proposed adaptive discretisation framework.

Minor comments

The minor comments mainly concern figure presentation, the description of data access services, additional references, and methodological clarifications aimed at improving reproducibility. These points have been carefully considered and will be addressed in the revised manuscript.

Several figure-related improvements will also be implemented. Panel labels will be repositioned where necessary to avoid overlap with graphical elements and map annotations, and figure quality will be improved through the use of higher-resolution or vector graphic outputs where appropriate.

An explicit reference to the European HF Radar Node will be added. In addition, the description of data accessibility will be expanded to clarify that datasets can be accessed through both THREDDS and ERDDAP services provided by the European HF Radar Node.

No additional threshold on the minimum number of radial measurements per site or on the total number of radial observations was applied. This clarification will be added to the manuscript to improve reproducibility.

These revisions are intended to improve clarity, transparency, and reproducibility, while leaving the methodology and scientific conclusions of the study unchanged.