https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

An adaptive unstructured grid for HF radar current mapping based on constrained k-means clustering

Bartolomeo Doronzo[1,2], Stefano Taddei[2] and Carlo Brandini[1,3]

[1]Istituto di Scienze Marine ISMAR-CNR, 50019 Firenze, Italy
[2]Consorzio LaMMA, 57126 Livorno, Italy
[3]Consorzio LaMMA, 50019 Firenze, Italy

Correspondence to: Bartolomeo Doronzo (bartolomeo.doronzo@cnr.it)

Abstract. High-Frequency (HF) radar systems provide high-resolution observations of surface currents in coastal oceans, but their effective representation strongly depends on the spatial discretization used to combine radial measurements. Standard regular grids impose a uniform resolution across the domain, despite the highly heterogeneous sampling geometry of HF radar systems, leading to inefficient use of observations and limited effective resolution in key coastal regions.

We propose a fully automated, non-uniform adaptive gridding framework specifically designed for HF radar applications. The method defines grid nodes directly from the spatial distribution of radar observations, redistributing resolution according to local data availability while preserving a consistent discretization of the domain. The resulting unstructured grid naturally refines resolution in densely sampled near-coastal areas and adopts coarser spacing offshore. Grid nodes can be associated with either Voronoi polygons or equivalent-area circular supports, providing flexibility in the geometrical representation of spatial support without altering node placement.

The proposed discretization yields a more realistic spatial representation of HF radar surface currents, reducing artefacts linked to uniform gridding and enhancing the visibility of coastal circulation structures. Comparison with surface drifter observations indicates that this improved spatial coherence is achieved without compromising overall reconstruction skill. The framework provides a radar-aware basis for current mapping, interpolation, and data assimilation applications where geometrical consistency and adaptive resolution are critical.

1 Introduction

High-Frequency (HF) radars are widely used to monitor surface currents in coastal regions, providing continuous, high-resolution, and synoptic observations over spatial scales ranging from a few kilometres to several tens of kilometres (Stewart and Joy, 1974; Paduan and Graber, 1997; Wyatt, 2014). Their capability to resolve rapidly evolving surface circulation patterns makes HF radar observations essential for a variety of operational and scientific applications, including search and rescue activities (Bellomo et al., 2015), pollution dispersion and oil-spill tracking (Abascal et al., 2009), ecosystem and biophysical studies (Kaplan et al., 2005), sediment transport analyses, and the validation of coastal ocean circulation models (Paduan et al., 2006; Tonani et al., 2019).

Despite their operational maturity, HF radar measurements are affected by several sources of uncertainty, including antenna-pattern distortions, geometric dilution of precision (GDOP), and the intrinsically anisotropic and heterogeneous sampling geometry of radar observations (Emery et al., 2004; Barrick, 2006; Lipa et al., 2006). These factors complicate the reconstruction of two-dimensional surface current vector fields from radial velocity components and introduce spatially varying uncertainty in the derived products.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

Most operational HF radar processing systems combine radial measurements on regular, structured grids using least-squares or related vector reconstruction techniques (Lipa and Barrick, 1983; Paduan and Cook, 1997). While computationally efficient and straightforward to implement, regular grids do not adequately represent the heterogeneous spatial distribution of HF radar observations. Due to antenna-beam spreading, geometric divergence, and range-dependent coverage limitations, radial measurements are generally sparse offshore and progressively denser toward the coast (Paduan and Washburn, 2013).

As a result, offshore grid cells often contain only a small number of radial measurements, leading to noisy or weakly constrained surface current estimates, whereas nearshore grid cells may aggregate a large number of observations, effectively oversmoothing small-scale coastal circulation features. This mismatch between fixed spatial discretization and variable data availability produces uneven vector accuracy, artificial spatial smoothing, and a loss of effective resolution in regions where fine-scale dynamics are most relevant (Cosoli et al., 2010; Cosoli and Bolzon, 2014). In poorly sampled areas, interpolation on regular grids can further introduce artefacts and physically unrealistic patterns, particularly near coastlines and land boundaries.

Adaptive and unstructured grids are widely employed in coastal ocean modeling to better resolve complex geometries, steep gradients, and multiscale dynamics (e.g. Fringer et al., 2006; Chen et al., 2006). In contrast, their application to HF radar data processing has been limited. Existing efforts to improve radar-derived surface currents have primarily focused on refining inversion methodologies, incorporating physical constraints, or enhancing error characterization. Examples include variational interpolation and physically constrained reconstruction schemes (e.g. Barth et al., 2021), signal-processing-based approaches such as the signal processing for flow mapping (SPFM) method (Shi et al., 2022), and the integration of HF radar observations with complementary in situ measurements, such as acoustic Doppler current profilers (ADCPs; Manso-Narvarte et al., 2020). While these approaches address aspects of vector reconstruction and uncertainty, they do not explicitly resolve the underlying mismatch between fixed spatial discretization and the heterogeneous sampling geometry of HF radar systems.

Several previous studies have explored alternative strategies to improve the reconstruction of HF radar surface currents, including unstructured discretisations, variational interpolation methods, and physically constrained inversion approaches. However, existing methodologies generally focus on the reconstruction or interpolation stage, while the spatial discretization itself is typically prescribed a priori and remains decoupled from the underlying radar sampling geometry. The intrinsic characteristics of HF radar observations naturally motivate discretization strategies that adapt spatial resolution to the local information content of the data. In this work, we introduce a non-uniform, adaptive gridding framework for HF radar applications that is explicitly data-driven and radar-aware. The proposed approach defines the spatial discretization directly from the distribution of radial observations, ensuring that grid resolution varies coherently with measurement density and geometrical constraints.

Within this framework, a constrained K-means clustering is employed as a practical means to determine the spatial distribution of grid nodes while enforcing balanced observational support across the domain. The resulting unstructured grid provides an information-adaptive discretization of the radar field, enhancing the representation of coastal circulation structures while maintaining statistical robustness in sparsely sampled regions. Additional processing steps, including outlier filtering, coastline-aware trimming of Voronoi polygons, and selection criteria based on radial density and geometric constraints, further improve the physical consistency of the derived surface current fields.

Validation against surface drifter trajectories, a commonly used benchmark for HF radar performance assessment (e.g. Kalampokis et al., 2016), demonstrates that the proposed adaptive discretization improves spatial coherence while maintaining pointwise error statistics comparable to those obtained using regular grids.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

2 Methodology

2.1 Dataset and study area

The study area is located in the northern Tyrrhenian Sea and includes the Tuscan coastline and adjacent offshore regions, extending from Elba Island across the Tuscan Archipelago to the Gulf of La Spezia. This region is characterized by complex coastal and oceanographic dynamics, shaped by insular and coastal morphology, highly variable bathymetry, and seasonally modulated atmospheric forcing. These factors contribute to strong spatial heterogeneity in surface circulation and make the area well suited for testing adaptive discretization strategies for HF radar observations.

The dataset analysed in this study is produced by the HF radar network operated by the LaMMA Consortium in cooperation with the Tuscany Region. The network consists of three CODAR SeaSonde direction-finding stations operating at a central frequency of 13.5 MHz: San Vincenzo (SVIN), Livorno (LIVO), and Isola del Tino (TINL). The SVIN and LIVO stations have been operational since 2015, while the TINL station has been active since 2018. Together, these sites provide coverage of a broad sector of the Ligurian–Tyrrhenian transition region, including the Tuscan Archipelago and the Gulf of La Spezia.

All radars operate with a bandwidth of 100 kHz, yielding an approximate radial range resolution of 1.5 km and an angular resolution of 5°. Radial surface current measurements refer to an effective sampling depth of approximately 0.9 m below the sea surface. Data are acquired at an hourly temporal resolution, and under favourable propagation conditions each station typically achieves a maximum range of 80–100 km.

Radial data from all stations are automatically transmitted to the central processing infrastructure of the European HF Radar Node, where they are archived and processed following standardized procedures (Roarty et al., 2019). Radial and total current products undergo quality control following established European HF radar guidelines, and final datasets are distributed in netCDF format through the THREDDS data servers of the European HF Radar Node (European High Frequency Radar Node, THREDDS Data Server, accessed on 24 April 2026).

2.2 Drifter Dataset

Surface drifter observations were used to validate HF radar currents and assess the impact of different gridding strategies. In October 2020, during the DDR20 field campaign in the southern Ligurian Sea, 26 drifters were deployed within the radar footprint to form an approximately 3 × 3 km array centred at 43°16′ N, 10°14′ E over waters approximately 100 m deep. For this experiment, CODE-type drifters (MAXO) and equipped with SPOT Globalstar TRACE modules, were used providing high-accuracy GPS positions (<10 m) at 5 min intervals.

CODE drifters sample currents in the upper meter of the water column and are minimally affected by wind forcing. Their wind slippage is typically 0.1 % of the wind speed, while wind-driven currents are of the order of 1 %. The resulting velocity accuracy is approximately 3 cm s^{-1}.

The present study focuses on the period between 8 October and 8 November 2020, corresponding to the temporal window of the drifter campaign. This ensures optimal temporal overlap between HF radar coverage and Lagrangian observations allowing a consistent comparison of different gridding strategies under highly variable flux conditions. For this reason, only the LaMMA radars active during the DDR20 experiment (SVIN, LIVO, and TINL) were included in the analysis. Further details on the campaign design, drifter behaviour, and radar–drifter comparisons are provided in previous studies (Doronzo et al., 2025a, b).

To complement the DDR20 dataset, trajectories from CARTHE-type drifters released by IFREMER near Bastia (Corsica) between May and September 2019 were also included. These instruments, sampled at 5 min intervals, follow currents

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

within the upper 0.6 m and exhibit minimal wave rectification and wind influence. Previous studies deploying paired CODE and CARTHE drifters have shown consistent behaviour, indicating that both platforms reliably track near-surface currents.

Additional details on the experimental drifter dataset, including full metadata, are available through the public data release of the DDR20 experiment (Doronzo et al., 2025c).

2.3 Grid Construction

The adaptive grid is designed to provide a spatial discretization of the radar observation domain that is consistent with the heterogeneous sampling geometry of HF radar systems. Instead of prescribing grid resolution a priori, the proposed approach allows resolution to vary across the domain in response to data availability, ensuring a closer alignment between spatial discretization and observational content.

A key principle is the separation between the definition of grid nodes and the geometrical representation of their spatial support. This design choice provides the flexibility required to handle complex coastal geometries and strongly non-uniform sampling patterns, while preserving a clear and consistent definition of the grid. The resulting adaptive discretization framework and its individual components are described in the following sections.

2.3.1 Overview

The adaptive grid is defined through a data-driven procedure that redistributes grid nodes according to the spatial density and geometry of HF radar observations. The underlying rationale is to align spatial resolution with the information content of the data, providing finer resolution in densely sampled near-coastal regions and progressively coarser resolution offshore, where radar coverage is inherently sparser.

To achieve this, grid node locations are derived from the spatial distribution of radar radial intersection points using a clustering-based strategy designed to ensure balanced observational support across the domain. This approach explicitly addresses the strong heterogeneity of HF radar sampling, preventing excessive aggregation of observations in dense regions while maintaining sufficient support where data availability is limited.

Once the grid nodes are defined, their associated spatial support is introduced through a geometrical partition that is independent of node placement. This separation between node definition and spatial support allows different strategies for associating radial observations with grid nodes to be explored without modifying the underlying discretization. Mathematical and algorithmic details are provided in Appendix A.

2.3.2 Classical K-means: mechanism and limitations

When spatial discretization does not account for the heterogeneous distribution of HF radar observations, grids defined with uniform spacing or unconstrained clustering tend to produce imbalanced representations of the radar field. In densely sampled regions, such as near the coastline or along directions of beam convergence, a large number of observations may be assigned to a limited set of grid nodes, effectively oversmoothing spatial variability. Conversely, in sparsely sampled offshore regions, grid nodes may be supported by only a few observations, resulting in weakly constrained or noisy estimates.

Clustering approaches that do not regulate the number of observations associated with each node naturally reproduce the spatial density of the input data. While this behaviour may be appropriate for some classification problems, it is poorly

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

suited for defining spatial discretisation intended to provide homogeneous statistical support for radar current reconstruction across the domain.

2.3.3 Balanced (Constrained) K-means

To overcome these limitations, we adopt a balanced clustering strategy that explicitly redistributes grid nodes as a function of observation density. By enforcing constraints on the number of radial intersections contributing to each node, the resulting discretization ensures that densely sampled regions are represented by multiple closely spaced nodes, while sparsely sampled offshore areas retain larger spatial supports.

This balanced formulation directly links node placement to the information content of the radar observations, rather than to an a priori prescribed grid geometry. As a result, the adaptive grid naturally reflects the intrinsic sampling characteristics of the HF radar system and provides a more uniform basis for subsequent vector reconstruction. The specific clustering algorithm employed to implement this constrained redistribution is described in Appendix A.

2.3.4 Spatial partitioning and definition of equivalent cells

The grid nodes obtained through the balanced clustering define the core spatial discretization of the domain. To associate a spatial support with each node, a Voronoi tessellation is constructed from the node locations, yielding a unique and non-overlapping partition of the domain that serves as the reference geometrical representation of the adaptive grid. The size and shape of the resulting cells directly reflect the adaptive distribution of grid nodes.

While the Voronoi tessellation provides a rigorous spatial partition, individual cells may become elongated or irregular, particularly in near-coastal regions or under anisotropic radar sampling geometries. To address this, the area of each Voronoi cell is used to define an equivalent circular support with the same area, centred on the corresponding grid node. This equivalent-area representation preserves the spatial scale of the adaptive grid while providing a more isotropic and numerically stable support for the association of radial observations.

In selected configurations, radial measurements are associated with grid nodes using these equivalent circular supports rather than the polygonal Voronoi cells. This choice affects only the geometrical definition of the spatial support and does not alter node placement or grid resolution. The equivalent-area formulation also provides a flexible metric for controlled expansion of the search area in regions of limited data availability, consistent with standard operational practices in HF radar processing (CODAR Ocean Sensors, Inc., 2016).

2.4 Workflow for Grid Construction and Post-Processing Filters

The adaptive grid is generated through a streamlined and fully automated workflow, schematically illustrated in Figure 1.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

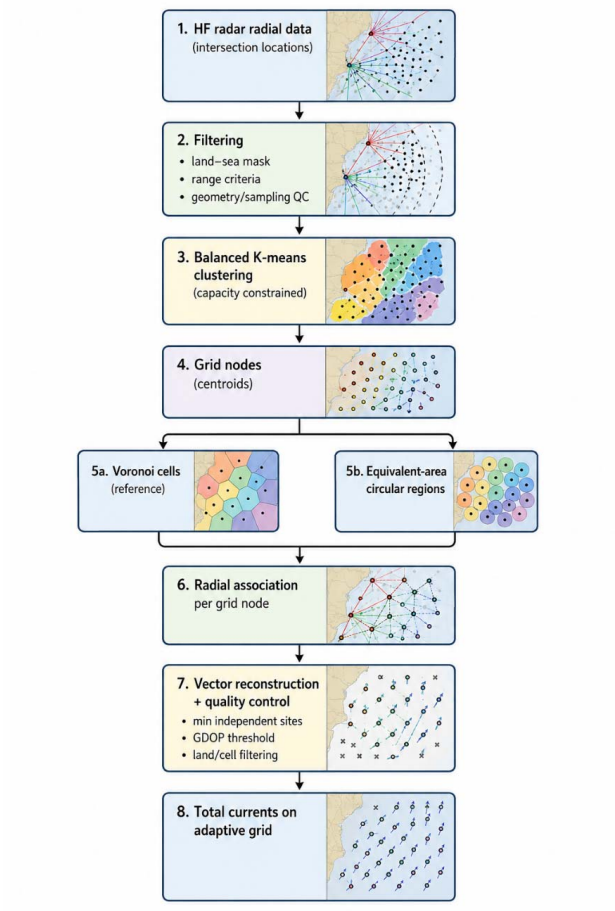

Figure 1. Schematic overview of the adaptive-grid workflow. Radar radial intersection locations are first filtered using land–sea masking and range- and geometry-based criteria. A balanced (capacity-constrained) K-means clustering is then applied to define grid nodes as centroids with comparable observational support. Spatial support is represented either by Voronoi polygons (reference configuration) or by equivalent-area circular regions (alternative configuration), followed by radial association, vector reconstruction, and quality control based on minimum independent radar contributions and GDOP thresholds. The workflow decouples node placement from the choice of spatial support, enabling controlled assessment of geometrical effects without modifying the underlying discretization.

The process begins with the filtering of radar radial observation locations using land–sea masking and range- and geometry-based criteria, in order to retain only points within the effective and physically consistent operational domain of the HF radar system. The filtered observation points are then used to define the spatial discretization of the domain through a balanced (capacity-constrained) K-means clustering. This step determines the positions of the grid nodes as centroids with comparable observational support, yielding a node distribution that adapts naturally to the heterogeneous HF radar sampling geometry, with higher density in near-coastal regions and progressively coarser spacing offshore. Once the grid nodes are defined, a geometrical representation of their spatial support is introduced to associate radial observations with each node. In the reference configuration, this association is based on a Voronoi tessellation derived from the centroid locations. In alternative configurations, radial observations are selected within equivalent-area circular regions centred on each centroid, allowing the impact of the geometrical definition of spatial support to be assessed without modifying the underlying node distribution. Finally, vector reconstruction and quality control are applied,

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

including the rejection of nodes intersecting land, the removal of anomalously large cells, minimum requirements on the number of independent radial contributions from distinct radar sites, and filtering based on geometric dilution of precision (GDOP). This workflow yields a flexible and radar-aware adaptive grid, in which node placement is decoupled from the choice of spatial support, providing a robust basis for the comparative analyses presented in the following sections.

2.5 Reconstruction and Validation of Surface Currents

Total surface current fields were reconstructed on both a regular grid and an adaptive unstructured grid using the same processing framework. In both configurations, radial velocity measurements from multiple HF radar sites were combined using a weighted least squares (WLS) approach. This formulation represents a generalization of the classical least squares method originally introduced for HF radar total current estimation (Lipa and Barrick, 1983), often referred to as unweighted least squares (UWLS) in standard applications (Kim et al., 2008).

For the regular grid configuration, radial observations were aggregated within fixed spatial cells, resulting in a uniform discretization of the domain. For the adaptive grid, radial measurements were instead associated with grid nodes using equivalent circular search regions derived from the area of the corresponding Voronoi cells. This approach provides a spatially varying definition of observational support while remaining conceptually consistent with the neighbourhood-based aggregation commonly adopted in operational HF radar processing.

To ensure robust geometric conditioning of the inversion, only grid nodes supported by radials from at least two distinct radar sites were retained. A geometric dilution of precision (GDOP) criterion was subsequently applied to identify and discard poorly constrained configurations characterized by unfavourable angular coverage. For each radar contributing to a given grid node, radial velocities were averaged and associated with a representative bearing, and a distance-based weighting scheme was adopted within the WLS inversion to account for the relative contribution of each site. For each grid configuration, both the number of contributing radials and the number of valid time steps were recorded to quantify spatial and temporal data coverage.

Following vector reconstruction, a post-processing quality control stage was applied to the total current fields. First, a spatial Hampel filter was used to identify and remove local outliers based on neighbouring vectors within a prescribed radius. This was followed by an angular standard deviation filter, which eliminates vectors whose direction deviates significantly from that of the surrounding flow. These filters were applied sequentially to reduce noise while preserving coherent circulation structures and improving the physical consistency of the reconstructed fields. In addition, all analyses were restricted to the effective radar coverage domain by applying an appropriate geographic mask.

Radar-derived surface currents were validated against independent surface drifter observations. For each radar time step, drifter positions within a ±30 min temporal window were selected to ensure consistency between Eulerian radar measurements and Lagrangian observations. Radar velocities were interpolated to the instantaneous drifter locations using a spatially weighted interpolation scheme, and only valid collocated pairs were retained after applying quality control filters to remove missing values and unrealistic velocities.

The GDOP value associated with each radar–drifter collocation was extracted from the nearest radar grid node. As a result, interpolated radar velocities may, in some cases, be associated with locally high GDOP values, reflecting the surrounding radar geometry rather than the conditioning of the interpolated solution itself.

The comparison between radar-derived and drifter velocities was performed using zonal and meridional components as well as current speed. Performance statistics included correlation coefficients and root mean square errors (RMSE), computed over the full set of valid collocations aggregated across the entire study period. In addition, robust linear regression based on the Theil–Sen estimator was used in scatterplot analyses to reduce sensitivity to outliers.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

Supplementary diagnostics, including the spatial separation between drifter positions and the nearest radar grid node and the temporal mismatch between observations, were also evaluated to characterize the representativeness of radar–drifter comparisons.

2.6 Diagnostic Checks for the Voronoi-Based Adaptive Grid

Following the construction of the adaptive grid, a set of diagnostic checks was applied to assess its internal consistency and suitability for total current reconstruction. These diagnostics are not intended as optimization criteria, but rather to verify that the proposed discretization achieves an appropriate balance between spatial resolution and observational support.

Basic geometrical properties of the grid were examined, including the consistency between centroid locations and their associated cells, the statistical distribution of Voronoi cell areas, and the spatial variability of radial density across the grid. These diagnostics are used to verify that the balanced redistribution of grid nodes effectively compensates for the offshore decrease in radar sampling density and avoids the formation of anomalously small or excessively large cells, particularly near complex coastlines and at the edges of the radar footprint.

Additional diagnostics focus on the availability of radar observations for total vector reconstruction. For each grid node, the number of contributing radials and the number of independent radar sites are evaluated prior to vector computation. Nodes that do not meet minimum requirements for reliable reconstruction are flagged and excluded from subsequent analyses. This allows the influence of grid resolution, radial association strategy, and radar geometry to be assessed independently of the reconstruction step itself.

The diagnostic results are used to guide the choice of key configuration parameters, such as the number of grid nodes and the strategy adopted for radial association, particularly in near-coastal regions where radar geometry strongly constrains data availability. These diagnostics also provide essential context for the comparative analyses presented in Sect. 3, especially when interpreting differences in spatial coverage and node utilization between regular and adaptive grids.

3 Results and Discussion

3.1 Grid Characteristics

Figure 2 illustrates the relationship between HF radar sampling geometry and the spatial distribution of grid nodes obtained through constrained K-means clustering. The radial observations acquired by the three radar sites (Fig. 2a) exhibit the characteristic HF radar sampling pattern, with high observation density near the radar locations and a progressive decrease offshore. The centroid distribution shown in Fig. 2b is derived from an idealised set of radar radial intersection points, representing the complete set of radials that could be generated in the absence of data gaps. Consistent with this ideal sampling geometry, the constrained clustering produces a higher density of centroids in near-radar regions and increasingly wider spacing offshore. This idealised centroid distribution provides the reference framework for the construction of the adaptive Voronoi grid.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

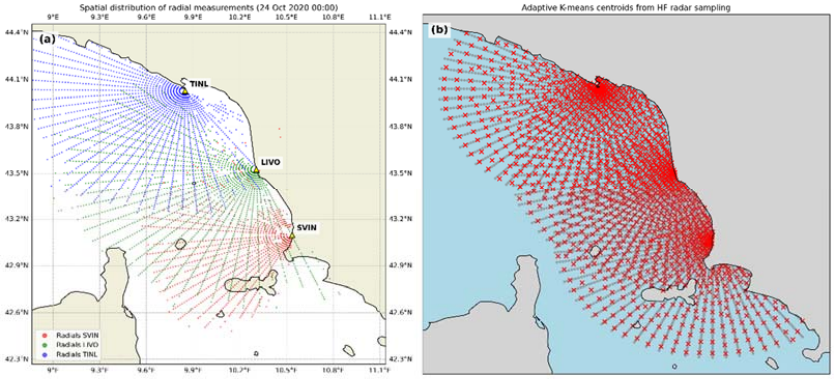

Figure 2. (a) Spatial distribution of HF radar radial measurements actually acquired by the three radar sites (San Vincenzo – SVIN, Livorno – LIVO, and Il Tino – TINL) during the selected time step. Radials are sampled at fixed azimuthal intervals (5°), resulting in a high density of observations close to each radar site and a progressive offshore decay of sampling density due to geometric and range limitations. (b) Spatial distribution of the 1300 centroids obtained from the constrained K-means clustering applied to an idealized set of HF radar radial sampling points, representing the complete set of radials that could be generated by the radar stations in the absence of data gaps and range limitations (the 1300-centroid configuration is shown here for illustrative and visualization purposes). Centroids cluster more densely in near-radar regions, reflecting the higher availability of potential radial measurements, while their spacing increases offshore in response to the reduced sampling density. This adaptive centroid distribution forms the basis for the subsequent construction of the Voronoi-based grid.

Building on this node distribution, Fig. 3a shows the corresponding Voronoi tessellation, with cells coloured according to the number of radial observations associated with each cell. The adaptive grid yields smaller cells in densely sampled near-coastal areas and larger cells offshore, resulting in a more homogeneous distribution of radial observations across the domain compared to a regular grid. Figure 3b presents the mean number of radials per Voronoi cell as a function of distance from the nearest radar site. For the configuration adopted in this study, the number of radials per cell remains approximately constant with increasing distance, typically around 4–5 observations per cell. This indicates that the constrained clustering effectively compensates for the natural offshore decrease in radar sampling density.

The number of radials associated with each Voronoi cell is primarily controlled by the number of centroids specified in the clustering procedure and can therefore be adjusted to tune the effective spatial resolution of the grid. In this study, the number of centroids was chosen to be comparable to the number of cells in the regular-grid configuration, enabling a direct comparison between the two approaches. Increasing the number of centroids enhances near-coastal resolution at the expense of fewer radials per cell, while the minimum achievable resolution remains constrained by the requirement of at least two independent radial observations per cell.

In addition to the polygonal Voronoi representation, each centroid can be associated with an equivalent circular region having the same area as the corresponding Voronoi cell (Fig. 3c). This equivalent-area formulation preserves the spatial distribution of grid nodes and modifies only the geometrical definition of the area used to associate radial observations. It is particularly advantageous in near-coastal regions, where Voronoi cells may become elongated or distorted due to radar geometry and coastline constraints. In such cases, equivalent circular regions provide a more isotropic search area and improve the robustness of radial selection and vector reconstruction. This formulation also allows moderate overlap between neighbouring cells under unfavourable radar geometries and is used in selected configurations in the analyses presented below.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

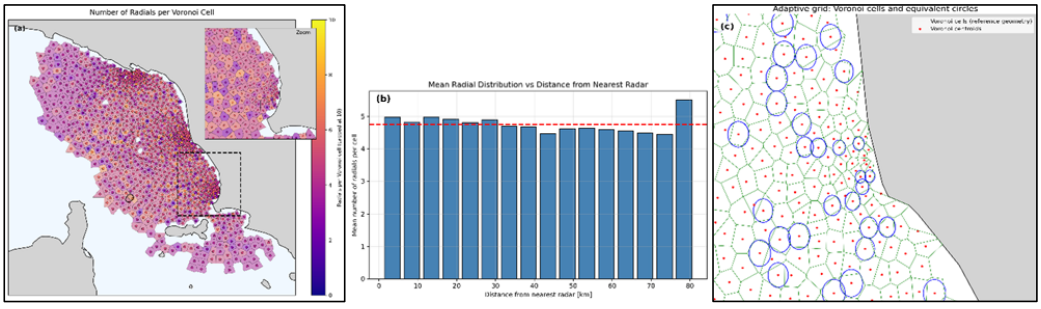

Figure 3. (a) Voronoi tessellation derived from the K-means centroids, coloured according to the number of radial measurements associated with each cell. The resulting adaptive, unstructured grid adjusts cell size to the local density of observations, with smaller cells in near-coastal, densely sampled regions and larger cells offshore. The dashed rectangle indicates a selected area that is shown in detail in the zoomed view on the right. (b) Mean number of radials per Voronoi cell as a function of distance from the nearest radar site. The average number of radials per cell remains approximately constant across the domain (typically around 4–5), reflecting the constraint imposed by the number of K-means centroids used in the clustering procedure. (c) Voronoi cells defining the adaptive grid together with their corresponding equivalent circular areas. Blue circles represent equivalent-area circles associated with each centroid, preserving the area of the corresponding Voronoi polygon while retaining the original spatial distribution of centroids. For clarity, only a subset of the equivalent circles is displayed.

Finally, Fig. 4a shows the distribution of Voronoi cell areas as a function of distance from the nearest radar site for the configuration with 1300 centroids. Within approximately 30 km from the radar sites, cell area generally increases with distance, reflecting the progressive reduction in sampling density. Beyond this range, cell sizes become more variable due to the overlap of radials from multiple radar sites. For comparison, a reference cell area of approximately 9 km² is commonly adopted in CODAR SeaSonde systems operating at similar frequencies. Relative to this reference scale, Voronoi cells are generally smaller near radar sites and progressively larger offshore, exceeding the regular-grid cell area beyond approximately 20–30 km. An example of a higher-resolution configuration based on 2300 centroids is shown in Fig. 4b, illustrating the finer spatial adaptation obtained by increasing the number of grid nodes.

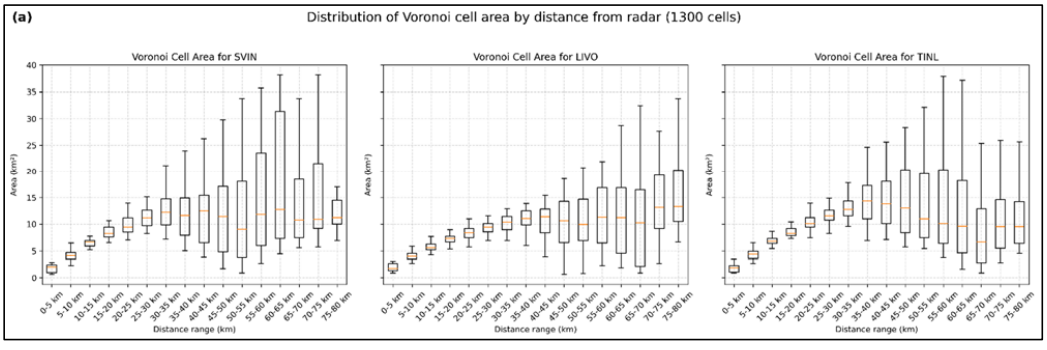

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

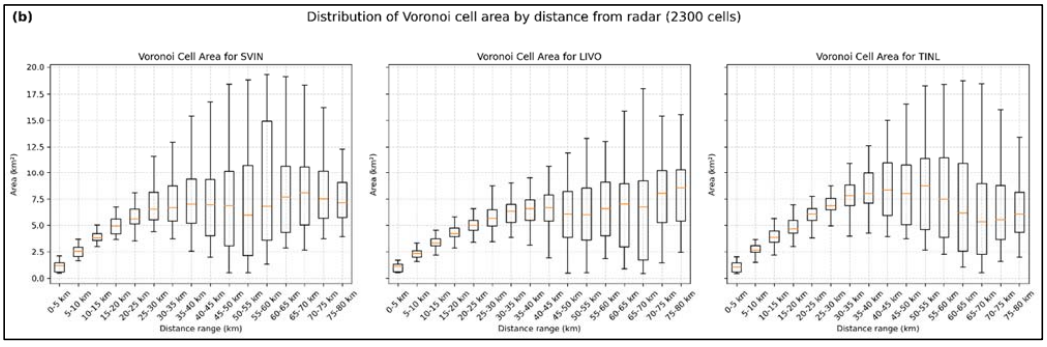

Figure 4. Boxplots of Voronoi cell area as a function of distance from the nearest radar site for two adaptive grid configurations: (a) 1300 cells and (b) 2300 cells. For each distance bin, the distribution of Voronoi cell areas is shown separately for the SVIN, LIVO, and TINL radar sectors. In both configurations, the median cell area and its variability increase progressively with distance from the radar, reflecting the adaptive nature of the Voronoi discretization, which allocates smaller cells in densely sampled near-range regions and larger cells in offshore areas. Beyond about 30 km, cell sizes exhibit greater variability due to the overlap of radials from multiple radar sites.

3.2 Comparison with Regular Grid

Fig. 5a illustrates the distribution of radar radial counts on a regular grid over the same domain. In contrast to the Voronoi-based unstructured grid, the regular grid exhibits a pronounced accumulation of radials in near-coastal cells, whereas offshore cells are often supported by only a few observations or remain sparsely populated. The corresponding histogram (Fig. 5b) confirms this behaviour, showing a rapid decrease in radial counts with increasing distance from the coast. This pattern reflects the intrinsic mismatch between the fixed spatial resolution of regular grids and the strongly heterogeneous sampling geometry of HF radar systems. As a result, regular grids lead to an unbalanced spatial distribution of observations, whereas the adaptive Voronoi approach provides a more homogeneous representation. The number of centroids can be adjusted to match the requirements of a given application, thereby controlling the trade-off between spatial resolution and statistical robustness. Fewer centroids produce larger cells with increased radial averaging, while a higher number of centroids yields smaller cells and finer spatial detail, provided that sufficient independent radial contributions are available.

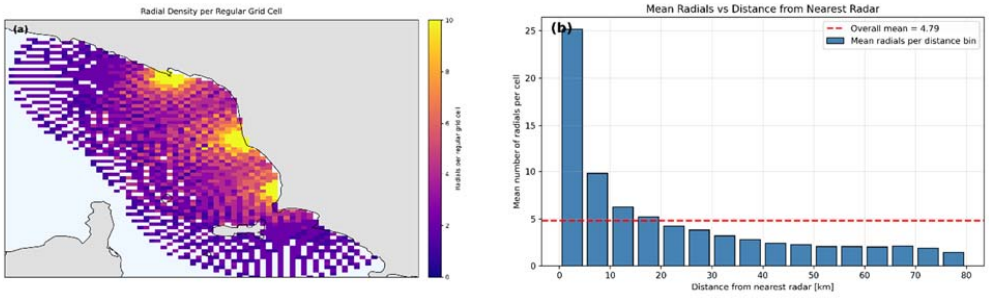

Figure 5. (a) Radial counts on a regular grid over the same HF radar domain. The fixed gridding leads to a pronounced aggregation of radial measurements within near-coastal cells, while offshore cells are often supported by only a few observations or remain poorly populated. This pattern illustrates the uneven spatial coverage characteristic of regular grids when applied to the intrinsically heterogeneous sampling geometry of HF radar systems. (b) Mean number of radials per regular grid cell as a function of distance from the nearest radar site. In contrast to the Voronoi-based adaptive grid (Figure 2b), the regular grid exhibits a strong decrease in radial counts with increasing distance, with very high values close to the coast and a rapid decay offshore. This distribution highlights the spatial imbalance inherent to fixed-resolution grids.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

Figure 6 (a–f) further illustrates the impact of the two gridding strategies by comparing the spatial distribution of grid-cell centroids effectively used for total current reconstruction in the vicinity of the SVIN, LIVO, and TINL radar sites. In the regular grid configuration, only a limited subset of nodes contributes to the reconstruction, with pronounced near-coastal clustering, particularly in the SVIN sector where multiple radials converge within the same fixed cells. While this configuration favours vector formation, it also leads to spatial aggregation and reduced effective resolution close to the coast.

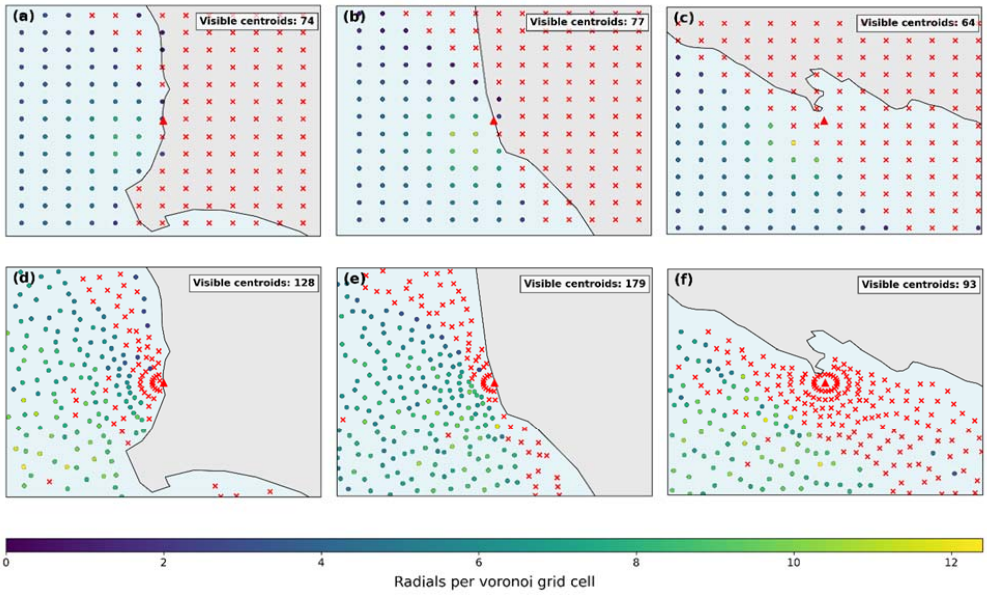

Figure 6. Spatial distribution of grid-cell centroids effectively used for total current reconstruction in the vicinity of the SVIN, LIVO, and TINL HF radar sites, comparing a regular grid discretization (panels a–c) and an adaptive Voronoi grid discretization (panels d–f). The regular grid (a–c) produces a limited number of effective centroids, with pronounced near-coastal clustering and a large fraction of excluded nodes. In contrast, the Voronoi grid (d–f) adaptively redistributes centroids according to radar sampling geometry and coastline configuration, increasing the number of effective grid cells, particularly in regions with favourable radial intersection geometry (e.g., the LIVO sector). The colour scale in panels (d–f) indicates the number of radials contributing to each Voronoi cell, while reduced centroid density elsewhere reflects geometric limitations of the observing system. Red asterisks mark grid nodes associated with invalid (NaN) values resulting from post-processing filters, including land masking, GDOP thresholds, and minimum independent-site requirements.

In contrast, the Voronoi-based grid redistributes centroids according to the underlying radar sampling geometry and coastline configuration, resulting in a substantially larger number of effectively used grid cells. This effect is particularly evident in regions characterised by favourable radial intersection geometry, such as the Livorno sector, where the adaptive discretisation increases spatial coverage and improves the utilisation of available observations. Conversely, in areas with limited angular diversity, some nearshore Voronoi centroids are excluded by the GDOP criterion, reflecting intrinsic geometric constraints of the observing system rather than limitations of the adaptive gridding approach. Overall, the Voronoi discretisation reallocates spatial resolution in a physically consistent manner, avoiding artificial overrepresentation in areas with insufficient data support.

These differences in grid structure and node utilisation should be considered when interpreting the radar–drifter comparisons presented in the following section, as variations in spatial representativeness and observational support may influence validation results independently of measurement accuracy. A complementary analysis was performed to examine the dependence of radar–drifter reconstruction error on radar geometry, quantified using the geometric dilution

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

of precision (GDOP). The results (Fig. 7a) show a clear and consistent increase in reconstruction error with increasing GDOP for both the regular and Voronoi-based grids, confirming the dominant role of radar geometry in controlling reconstruction quality. The two gridding approaches exhibit similar RMSE–GDOP relationships, indicating a comparable sensitivity to geometric degradation. Differences between the configurations are small and do not indicate a systematic improvement in pointwise reconstruction accuracy for the Voronoi-based grid. As expected, the number of radar–drifter collocations decreases with increasing GDOP (Fig. 7b), implying that higher GDOP classes are based on fewer observations and should therefore be interpreted with caution. In this context, the slight flattening or decrease of RMSE observed for the Voronoi-based grid at intermediate to high GDOP values is consistent with a more effective use of available observations under suboptimal geometric conditions, but it should not be interpreted as a robust reduction of reconstruction error. Similarly, the Voronoi-based grid tends to retain a comparable or slightly larger number of valid radar–drifter pairs at intermediate and higher GDOP values, indicating a marginal improvement in data availability under degraded geometry. However, this effect remains secondary compared to the dominant influence of GDOP. Overall, these results indicate that the main advantage of the Voronoi-based adaptive grid lies not in reducing reconstruction error under degraded geometry, but in providing a more balanced use of available observations, consistent with the spatial coverage differences discussed above.

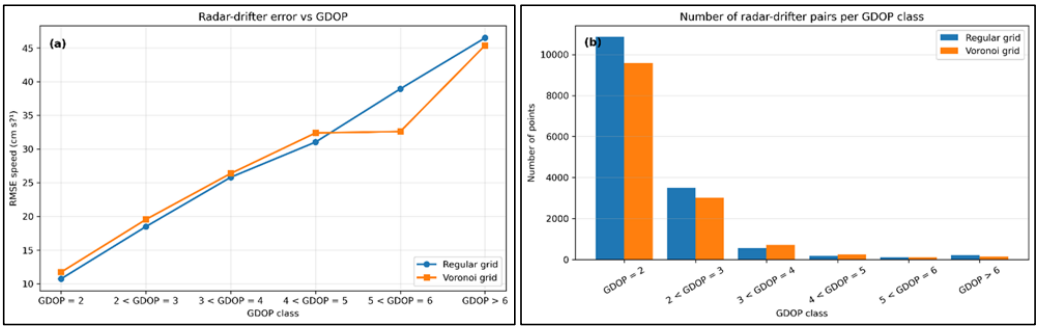

Figure 7. Dependence of radar–drifter current reconstruction error on radar geometry. (a) Root mean square error (RMSE) of radar–drifter current speed differences as a function of GDOP classes, shown for both the regular grid and the Voronoi-based adaptive grid. (b) Number of radar–drifter collocations within each GDOP class for the two grid configurations. The progressive decrease in sample size at higher GDOP values should be taken into account when interpreting the corresponding RMSE estimates.

3.3 Spatial representativeness and regional analysis of the radar–drifter comparison

The spatial characteristics of the radar-derived surface current fields were analysed by comparing maps reconstructed using a regular grid and an adaptive Voronoi-based grid. The comparison focuses on spatial coverage, continuity, and organisation of the reconstructed velocity fields, rather than on pointwise accuracy. Results are presented for a restricted sub-domain (Fig. 8a–f), allowing systematic differences between the two gridding strategies to be examined at the local scale. In the restricted area (Fig. 8), the regular grid reconstructions (panels a, c, e) exhibit a uniform but relatively sparse spatial sampling. The resulting velocity fields show noticeable gaps and reduced spatial continuity, with flow patterns that appear fragmented, particularly in regions affected by unfavourable radar geometry. These features are consistent across the analysed time steps and lead to a partial representation of mesoscale structures.

In contrast, the Voronoi-based reconstructions (Fig. 8b, d, f) provide a denser and spatially adaptive coverage of the same domain. The velocity fields exhibit improved spatial continuity and a more homogeneous filling of the area, allowing

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

mesoscale features such as rotational structures and shear zones to emerge more clearly and with greater dynamical consistency. Compared to the regular grid, these structures appear more coherent and less affected by sampling gaps.

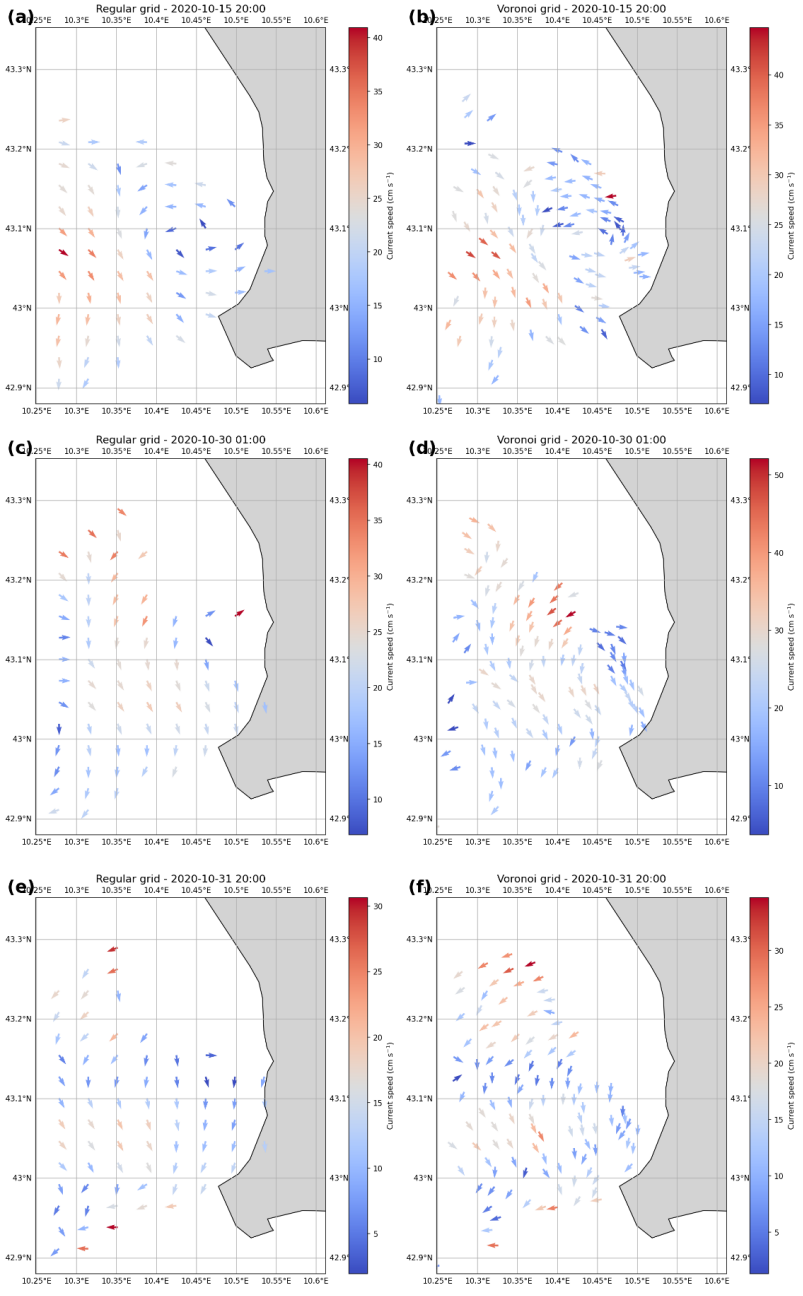

Figure 8. Comparison of radar-derived surface current fields reconstructed over a restricted sub-domain using a regular grid and an adaptive Voronoi-based grid. Panels (a, c, e) show reconstructions obtained with the regular grid, while panels (b, d, f) show the corresponding Voronoi-based reconstructions for three representative time steps. In the regular-grid cases (a, c, e), the velocity fields exhibit a relatively sparse spatial sampling, with noticeable gaps and reduced spatial continuity, particularly in areas affected by less favourable radar geometry. Flow structures appear fragmented and, in some instances, only partially

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

resolved. In contrast, the Voronoi-based reconstructions (b, d, f) provide a denser and spatially adaptive representation of the flow, resulting in improved coverage and enhanced spatial coherence. Mesoscale features such as rotational patterns and shear zones are more clearly expressed and dynamically consistent, reflecting the ability of the adaptive grid to redistribute spatial support according to the local availability of radial measurements. In all panels, vectors are colour-coded by current speed (cm s^{-1}), and land areas are masked.

The differences observed at the local scale are further confirmed when considering the full radar footprint (Fig. 9).

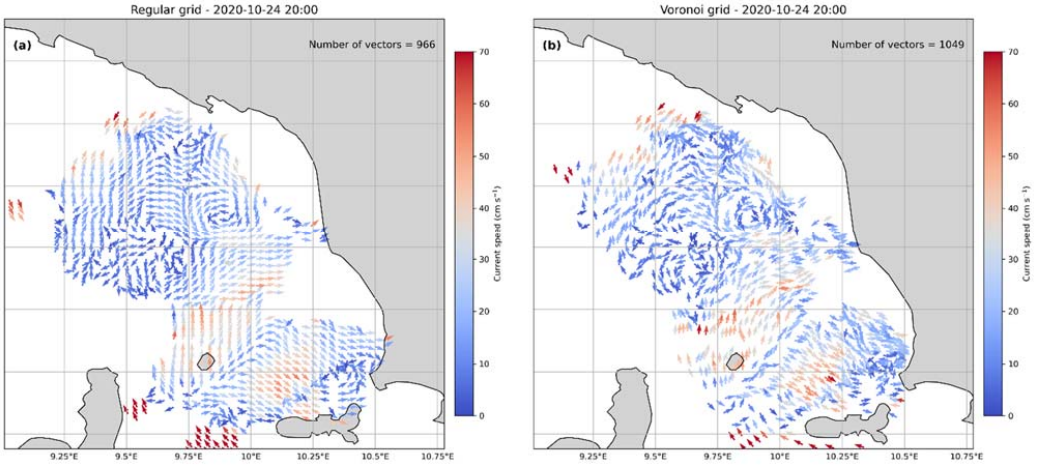

Figure 9. Surface current fields reconstructed over the full radar domain using a regular grid (a) and a Voronoi-based adaptive grid (b) for a representative time step. At the scale of the entire radar footprint, the Voronoi grid yields a more continuous and spatially complete field, particularly near the domain boundaries and in regions of heterogeneous sampling, while the regular grid shows increased fragmentation. This domain-scale comparison confirms that the differences in spatial coverage and coherence observed in the restricted sub-domain (Figure 8) persist at larger spatial scales. Vectors are coloured by current speed (cm s^{-1}).

At this scale, the Voronoi-based grid produces a more complete and continuous reconstruction, particularly near the edges of the radar coverage and along the coastline, where data availability is more heterogeneous. In contrast, the regular grid, while preserving a fixed spatial structure, results in sparser coverage and increased fragmentation of the velocity field in these areas. These differences reflect the distinct spatial support of the two approaches. The regular grid relies on a fixed and uniform geometry, whereas the Voronoi-based method adapts the spatial discretisation to the local distribution of radial measurements. This adaptive behaviour enables a more effective use of the same input data without modifying the applied quality-control criteria and reduces interpolation artefacts, such as unrealistic velocity vectors extending over land or across poorly sampled regions. As a result, under complex geometric conditions, the Voronoi-based grid better preserves spatial continuity and the representation of dynamically consistent surface current patterns.

Overall, these results are consistent with those presented in Sect. 3.2. The Voronoi-based approach provides a more robust representation of the flow field in terms of spatial coverage and coherence. This improved representativeness is particularly relevant for applications involving the identification of transport pathways and mesoscale dynamics, where an accurate reconstruction of spatial structures is as important as the limitation of local errors.

To further quantify the impact of the two gridding strategies on radar–drifter comparisons, the spatial separation between drifter positions and their nearest radar reconstruction nodes was analysed. Figure 10 shows the distribution of these distances over the full domain for the regular grid (Fig. 10a) and the Voronoi-based grid (Fig. 10b).

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

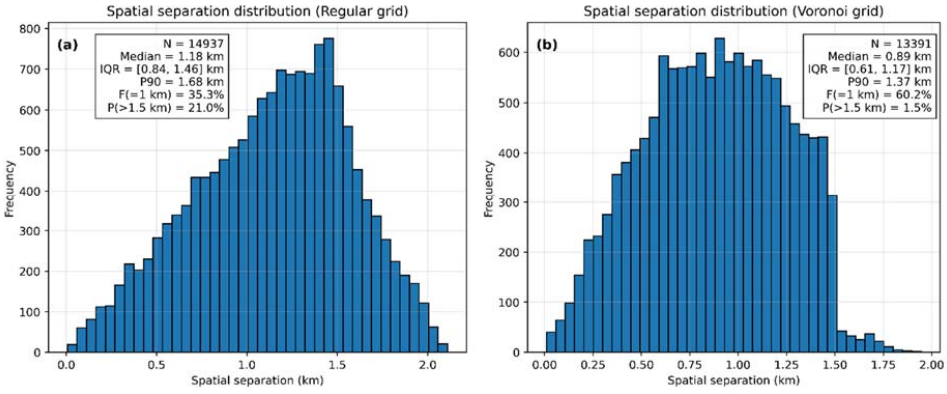

Figure 10. Distributions of the spatial separation between drifter positions and their nearest radar reconstruction nodes over the full radar domain, used as a proxy for the spatial representativeness of radar–drifter collocations. The left panel (a) shows results for the regular grid, while the right panel (b) refers to the Voronoi-based adaptive grid. Compared to the regular grid, the Voronoi-based approach yields systematically shorter radar–drifter separations, with a higher fraction of collocations within 1 km and a strongly reduced upper tail at larger distances. These differences reflect the distinct spatial support of the two gridding strategies and their ability to place reconstruction nodes closer to drifter trajectories in regions of heterogeneous radar sampling.

A clear difference emerges between the two configurations. The regular grid exhibits a distance distribution shifted towards larger separations, with a median radar–drifter distance of 1.18 km and an interquartile range of 0.84–1.46 km. Consequently, only about 35% of drifter positions are located within 1 km of the nearest grid node, while approximately 21% of collocations occur at distances greater than 1.5 km. In contrast, the Voronoi-based grid yields systematically shorter separations. The median distance decreases to 0.89 km and the interquartile range narrows to 0.61–1.17 km, indicating a more concentrated distribution. The fraction of close collocations increases substantially, with more than 60% of drifter positions lying within 1 km of the nearest Voronoi centroid. At the same time, the upper tail of the distribution is reduced: the 90th percentile decreases from 1.68 km for the regular grid to 1.37 km, and the proportion of separations exceeding 1.5 km drops to less than 2%.

These differences reflect the distinct spatial organisation of the reconstruction nodes generated by the two gridding approaches. The fixed geometry of the regular grid imposes a uniform spatial sampling that is independent of the underlying data density, leading to larger separations in regions with sparse radar coverage. Conversely, the Voronoi-based approach adapts the spatial distribution of nodes to the radar sampling geometry, resulting in a closer spatial correspondence between radar-derived currents and in situ observations. Such differences have direct implications for radar–drifter validation, as shorter spatial separations are generally associated with reduced representativeness errors and a lower impact of spatial interpolation mismatch. In line with this behaviour, the Voronoi-based grid improves the spatial consistency of radar–drifter comparisons.

This interpretation is supported by the scatter observed in Fig. 11, where the Voronoi-based reconstructions (right panels) exhibit a more compact relationship between radar and drifter velocity components compared to the regular-grid case (left panels). In particular, the reduced spatial separation is reflected in a lower dispersion around the fitted relationships, especially in regions characterised by complex coastal geometry, such as the spatially filtered area near the Livorno radar site. These results indicate that part of the observed scatter in radar–drifter comparisons can be attributed to spatial representativeness effects associated with the underlying gridding strategy, rather than solely to measurement noise or temporal mismatch.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

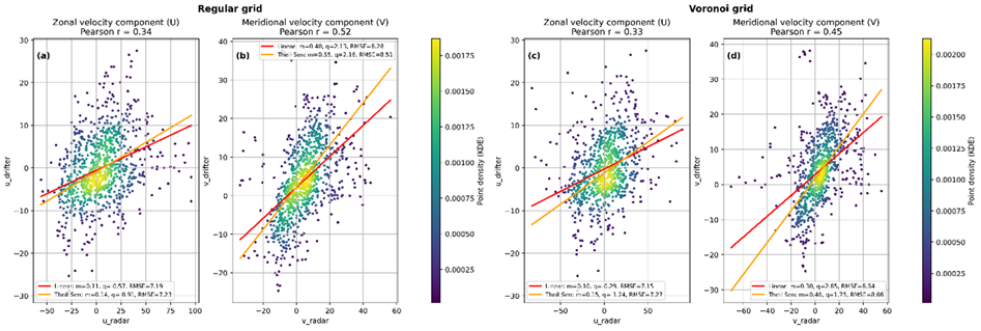

Figure 11. Scatterplot comparison between radar-derived and drifter-derived surface current velocities in a spatially filtered sub-domain near the Livorno (LIVO) HF radar site. Panels (a, b) show the zonal (U) and meridional (V) components for the regular grid, while panels (c, d) show the corresponding results for the Voronoi-based adaptive grid. Only radar–drifter collocations within this high-quality coverage area are considered. Point colours indicate data density (KDE), with Theil–Sen robust linear fits overlaid. Pearson correlation coefficients and RMSE values are reported.

Despite these differences, both grids exhibit similar statistical performance in terms of RMSE and correlation, indicating that pointwise accuracy is only weakly sensitive to the gridding approach. This suggests that representativeness effects primarily influence the spatial structure of the reconstructed fields and the dispersion of the comparison, while the dominant source of variability remains associated with radar geometry and spatial sampling.

The Voronoi-based grid provides a more spatially consistent representation of the flow field, particularly under geometrically complex conditions, by aligning the discretisation with the underlying radar sampling geometry. This spatially stratified evaluation enables an assessment of radar performance under both typical and constrained geometric conditions. Importantly, the same spatial filtering criteria and validation procedure were applied to both the regular and Voronoi-based fields, ensuring a direct and objective comparison across spatial scales.

3.4 Robustness and data availability

In addition to the physical robustness discussed in the previous sections, the Voronoi-based adaptive grid provides relevant advantages in terms of data representativeness, availability, and computational efficiency. By concentrating nodes in regions with dense and geometrically well-conditioned observational support, while progressively reducing spatial resolution in sparsely sampled offshore or poorly constrained areas, the Voronoi discretization adapts naturally to the intrinsic information content of the radar measurements. This adaptive behaviour leads to a more efficient use of the available observational information: grid nodes are allocated primarily where the radar system effectively constrains the surface velocity field, whereas regions characterized by limited coverage, high GDOP, or weak intersection geometry contribute a reduced number of reconstruction points. As a consequence, the Voronoi formulation reduces the number of reconstructed vectors compared to a regular grid spanning the same geographical domain, without altering the underlying physical reconstruction scheme or relaxing quality-control criteria.

From a practical standpoint, this reduction in effective spatial degrees of freedom directly translates into lower storage requirements for both intermediate products (e.g., radial aggregations and WLS contributions) and final gridded fields, as well as decreased computational cost for vector reconstruction, filtering procedures, and statistical post-processing. In the present application, this effect is reflected in a substantial reduction of stored dataset size (approximately a factor of five compared to the regular grid), despite comparable statistical performance in terms of correlation and RMSE. Importantly, these computational and storage benefits should not be interpreted as a simple compression strategy, but rather as a consequence of improved spatial sampling efficiency with respect to the radar observational geometry. In this

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

sense, the Voronoi grid effectively acts as an information-aware representation of the velocity field, avoiding the inclusion of low-information or weakly constrained regions that are still present in a uniform discretization. This becomes particularly relevant in long-term reconstructions, multi-year climatologies, and operational systems, where both scalability and data manageability are critical constraints.

4. Conclusions

This study demonstrates that an unstructured, radar-aware gridding strategy based on constrained k-means clustering and Voronoi tessellation provides a consistent and information-adaptive representation of HF radar surface current observations. Rather than constituting an incremental improvement over conventional regular grids, the proposed approach introduces a shift in perspective, from a fixed and uniform discretisation to a geometry-aware framework in which spatial resolution reflects the intrinsic distribution of observational information.

By explicitly accounting for the radar sampling geometry, the method redistributes grid nodes and spatial support according to the local data content, effectively aligning the discretisation with the degrees of freedom of the observing system. This results in an improved representation of spatial coverage and flow organisation across the radar domain, with enhanced utilisation of observational data under heterogeneous sampling conditions.

Validation against drifter observations indicates that, while standard pointwise metrics such as RMSE and correlation remain comparable between regular and adaptive grids, the Voronoi-based approach better captures the spatial structure of the flow field. In particular, the results show that differences between radar and drifter observations are partly attributable to spatial representativeness effects, which are reduced when using an adaptive discretisation that more closely follows the radar geometry. The main advantage of the Voronoi-based grid therefore lies in its ability to provide a more coherent and physically consistent spatial representation of surface currents. Near the coast, where radar observations are dense and well-conditioned, the adaptive grid resolves finer-scale features, while offshore and in geometrically constrained regions it reduces spatial resolution in accordance with the available data support. This leads to a more balanced distribution of effective reconstruction nodes, improving continuity and reducing artefacts associated with fixed-grid interpolation. More generally, the proposed approach can be interpreted as an information-aware discretisation of the velocity field, in contrast to conventional grid-based representations that do not explicitly account for the spatial variability of observational support. This perspective is particularly relevant for HF radar systems, where sampling is inherently observation-driven and strongly controlled by geometry.

In addition to its representational advantages, the adaptive framework provides practical benefits in terms of computational efficiency and data management. By limiting the number of reconstruction nodes in poorly constrained regions, the method reduces storage requirements and computational costs without altering the underlying reconstruction scheme or quality-control criteria. These features are especially relevant for large-scale applications, including long-term reconstructions, multi-year climatologies, and operational monitoring systems.

The adaptive, radar-aware gridding framework developed in this study also opens perspectives for future applications. Its compatibility with irregular observational geometries makes it well suited for complex coastal configurations and multi-radar networks, where conventional grids are often inefficient. Furthermore, the alignment between discretisation and observational structure offers a natural interface for integration with numerical models and data assimilation systems, where representativeness errors play a critical role.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

Future work may explore the extension of this approach to incorporate uncertainty-weighted reconstructions, temporal adaptivity, and hybrid gridding strategies, further enhancing the exploitation of HF radar observations under variable sampling conditions.

Appendix A: Voronoi tessellation and constrained (balanced) K-means clustering

To construct a spatially coherent grid adapted to the heterogeneous sampling geometry of HF radar systems, we adopt a two-step approach consisting of (i) constrained (balanced) K-means clustering to define grid node locations and (ii) a Voronoi-based geometrical partition to represent the spatial support associated with each node. This formulation addresses the limitations of classical clustering methods when applied to strongly non-uniform geospatial datasets such as radar-derived radial observation points.

A1. Classical K-means: formulation and limitations

Let $X = \{x_i, \dots, x_N\} \subset \mathbb{R}^2$, be a set of spatial data points, where x_i denotes the Cartesian coordinates of a radar radial intersection. Classical K-means partitions X into Kdisjoint clusters $C_i, \dots, C_K$ by minimizing the within-cluster sum of squared distances:

$$\min_{C_1,\dots,C_K} \sum_{k=1}^{k} \sum_{x_i \in C_k} \|x_i - \mu_k\|^2 , \quad (1)$$

where

$$\mu_k = \frac{1}{|C_k|} \sum_{x_i \in C_k} x_i \quad (2)$$

is the centroid of cluster C_K.
The solution is typically obtained using Lloyd's algorithm (Lloyd, 1982), which alternates between an assignment step,

$$x_i \in C_k \ \ if \ \ k = arg \min_j \|x_i - \mu_k\|^2 \quad (3)$$

and an update step:

$$\mu_k \leftarrow \frac{1}{|C_k|} \sum_{x_i \in C_k} x_i, \quad (4)$$

until cluster assignments no longer change.
Despite its simplicity and computational efficiency, classical K-means presents two major limitations in the context of HF radar data:

1. Lack of control on cluster size.
   The objective function contains no explicit term regulating the cardinality $|C_k|$ of the clusters. When the spatial distribution of points is highly heterogeneous, as in HF radar sampling, this leads to strongly unbalanced clusters, with excessive aggregation in densely sampled near-coastal regions and insufficient support offshore.

2. Premature convergence.
   Lloyd's algorithm is prone to convergence toward local minima where assignments stabilize even if cluster size imbalance persists, a well-known consequence of the piecewise-constant nature of the assignment step.

These limitations are documented in seminal work (MacQueen, 1967; Lloyd, 1982) and in subsequent analyses of clustering imbalance in spatial datasets.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

A2. Balanced (capacity-constrained) K-means

To overcome these issues, we employ a capacity-constrained, or balanced, K-means formulation in which each cluster is required to contain a number of points within prescribed bounds. The constrained optimization problem is written as

$$\min_{C_1,\ldots,C_K} \sum_{k=1}^{k} \sum_{x_i \in C_k} \|x_i - \mu_k\|^2 \text{ subject to } L \leq |C_K| \leq U, \forall k, \quad (5)$$

where L and U are lower and upper bounds on cluster cardinality, estimated from the ratio N/K.

In this formulation, the assignment step is modified to satisfy the capacity constraints, while centroid updates retain the same expression as in the classical algorithm. Iterations are repeated until cluster memberships stabilize or a prescribed stopping criterion is reached. The precise implementation details follow standard approaches proposed for constrained clustering in spatial and optimization problems (Bradley and Fayyad, 1998; Bradley et al., 2000; Malinen and Fränti, 2014).

Imposing capacity constraints ensures a more homogeneous distribution of cluster populations and therefore of centroids across the domain. In the HF radar context, this prevents oversampling of regions characterized by dense radial intersections and stabilizes clusters in sparsely sampled offshore areas. As a result, centroids are redistributed according to both the spatial distribution of observations and the requirement of balanced informational support, providing a physically meaningful adaptive discretization of the radar field.

A3. Voronoi tessellation

The centroids obtained from the constrained K-means clustering define the generating sites for a Voronoi tessellation, which provides a geometrical representation of the spatial support associated with each grid node (Aurenhammer, 1991). Given the set of centroid locations $\{\boldsymbol{p_i}, \ldots, \boldsymbol{p_N}\} \subset \mathbb{R}^2$, the Voronoi region associated with centroid $\boldsymbol{p_k}$ is defined as:

$$V_k = \{x \in \mathbb{R}^2 : \|x - p_k\| \leq \|x - p_j\|, \forall_j \neq k\}. \quad (6)$$

The resulting Voronoi polygons form a non-overlapping partition of the domain, with cell size and shape reflecting the adaptive distribution of the centroids. The combined use of balanced clustering and Voronoi discretization yields a grid that respects radar geometry, mitigates biases induced by point-density gradients, and avoids unrealistic aggregation in regions of radial convergence.

A4. Equivalent-area representation of Voronoi cells

While Voronoi tessellations provide a unique spatial partition, the corresponding polygons may become elongated or irregular, particularly near coastlines or under anisotropic sampling geometries typical of HF radar systems. To provide a more isotropic representation of spatial support, the area A_k of each Voronoi cell is used to define an equivalent circular region with radius

$$r_{eq,k} = \sqrt{\frac{A_k}{\pi}} \quad (7)$$

The equivalent-area circular cell is centered at the corresponding centroid p_k and preserves the spatial scale of the Voronoi polygon while replacing its potentially complex geometry with a simple isotropic support. This transformation does not modify the distribution of grid nodes nor the adaptive resolution of the grid, which remain entirely determined by the constrained K-means clustering.

In the present study, the Voronoi tessellation is retained as the reference geometrical partition, while the equivalent-radius formulation is used in selected configurations to improve the robustness of radial association for vector reconstruction. This approach is particularly advantageous under complex coastal geometries or unfavorable radar configurations, and is

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

consistent with standard CODAR SeaSonde processing practices that combine radial velocities within circular search regions of prescribed radius (CODAR Ocean Sensors, Inc., 2016).

A5. Geometrical considerations

All distances, Voronoi constructions, and polygon areas are computed in a local planar projection to ensure metric consistency in centroid definition, distance evaluation, and area calculation.

Code and data availability

Code availability. The code used to generate the adaptive K-means/Voronoi grid from HF radar sampling geometry will be made publicly available upon acceptance of the manuscript.

During the review process, the code can be made available to editors and reviewers upon request, either as a compressed archive or via access-restricted repositories, in order to support reproducibility.

After acceptance, the code will be released through a public repository and archived in a trusted long-term repository with a persistent identifier (e.g. DOI).

Data availability. The present study uses HF radar NetCDF files that are produced and distributed by operational radar networks and are subject to the data policies of the respective providers. These data are not redistributed with this manuscript. All products generated within this study (sampling geometry files and adaptive grid outputs) can be fully reproduced using the released code and the original HF radar geometry files.

Author contributions

All authors contributed to the conception of the study. BD developed the methodology, implemented the code, and performed the analysis. BD prepared the manuscript. CB contributed to funding acquisition, provided supervision, and contributed to the review and editing of the manuscript. All authors reviewed and approved the final version of the manuscript.

Competing interests

The authors declare no conflict of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Acknowledgments

The radar data cited in this article are part of the observation network of the LaMMA Consortium, consisting of a multi-site coastal HF radar system operating in the northern Tyrrhenian and Ligurian Seas, developed within several regional and cross-border monitoring initiatives.

The drifter campaign is partly funded by the IMPACT (IMpact of Ports on marine protected areas: Cooperative Cross-border Actions) project, included in the Interreg Italia–Francia Marittimo 2014–2020 programme.

Artificial intelligence tools were used to support code implementation and debugging, and to improve readability and language in some parts of this paper. After using these tools, the authors reviewed and edited the content as needed. All scientific decisions, methodological choices, results, and final implementations were designed, validated, and approved by the authors, who take full responsibility for the content of the publication.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

References

Abascal, A. J., Castanedo, S., Medina, R., Liste, M., and Fernández, F.: Application of HF radar currents to oil spill modelling, *Mar. Pollut. Bull.*, **58**, 238–248, https://doi.org/10.1016/j.marpolbul.2008.09.020, 2009.

Aurenhammer, F.: Voronoi diagrams: a survey of a fundamental geometric data structure, *ACM Comput. Surv.*, **23**, 345–405, 1991.

Barrick, D. E.: Geometric dilution of precision in HF radar surface current mapping, *IEEE J. Oceanic Eng.*, **31**, 843–849, 2006.

Barth, A., Troupin, C., Reyes, E., Beckers, J.-M., and Alvera-Azcárate, A.: Variational interpolation of high-frequency radar surface currents using DIVAnd, *Ocean Dyn.*, **71**, 1287–1303, https://doi.org/10.1007/s10236-020-01432-x, 2021.

Bellomo, L., Griffa, A., Cosoli, S., Molcard, A., Farneti, R., Poulain, P.-M., Menna, M., and Falco, P.: Toward an integrated HF radar network in the Mediterranean Sea to improve search and rescue and oil spill response: the TOSCA project experience, *J. Oper. Oceanogr.*, **8**, 95–107, https://doi.org/10.1080/1755876X.2015.1087184, 2015.

Bradley, P. S. and Fayyad, U. M.: Refining initial points for K-means clustering, *Proc. ICML*, 1998.

Bradley, P. S., Bennett, K. P., and Demiriz, A.: Constrained K-means clustering, Microsoft Research Technical Report, 2000.

Malinen, M. I. and Fränti, P.: Balanced K-means for clustering, *LNCS*, **8621**, 32–41, https://doi.org/10.1007/978-3-662-44415-3_4, 2014.

Chen, C., Beardsley, R. C., and Cowles, G.: An unstructured grid, finite-volume coastal ocean model (FVCOM) system, *Oceanography*, **19**, 78–89, https://doi.org/10.5670/oceanog.2006.92, 2006.

CODAR Ocean Sensors, Inc.: *SeaSonde radial and total data file formats (LLUV)*, technical documentation, release 6, available at: http://support.codar.com/Technicians_Information_Page_for_SeaSondes/Manuals_Documentation_Release_8/File_Formats/File_LLUV.pdf, last access: 29 April 2026, 2016.

Cosoli, S., Mazzoldi, A., and Gačić, M.: Validation of surface current measurements in the northern Adriatic Sea from high-frequency radars, *J. Atmos. Oceanic Technol.*, **27**, 908–919, https://doi.org/10.1175/2009JTECHO680.1, 2010.

Cosoli, S. and Bolzon, G.: Accuracy of HF radar surface current mapping, *J. Atmos. Oceanic Technol.*, **31**, 876–895, 2014.

Doronzo, B., Bendoni, M., Taddei, S., Boccacci, A., and Brandini, C.: Validating HF radar current accuracy via Lagrangian measurements and radar-to-radar comparisons in highly variable surface currents, *Remote Sens.*, **17**, 1243, https://doi.org/10.3390/rs17071243, 2025a.

Doronzo, B., Bendoni, M., Taddei, S., Boccacci, A., and Brandini, C.: A combined HF radar and drifter dataset for analysis of highly variable surface currents, *Data*, **10**, 115, https://doi.org/10.3390/data10070115, 2025b.

Doronzo, B., Poulain, P.-M., Taddei, S., Boccacci, A., and Brandini, C.: CODE drifter observations from the DDR20 experiment in the NW Mediterranean Sea, *SEANOE*, https://doi.org/10.17882/105936, 2025c.

Emery, B. M., Washburn, L., and Harlan, J. A.: Evaluating radial current measurements from CODAR HF radars with moored current meters, *J. Atmos. Oceanic Technol.*, **21**, 1259–1271, 2004.

Fringer, O. B., Gerritsen, M., and Street, R. L.: An unstructured-grid, finite-volume, nonhydrostatic, parallel coastal ocean simulator, *Ocean Model.*, **14**, 139–173, https://doi.org/10.1016/j.ocemod.2006.03.006, 2006.

Kalampokis, A., Uttieri, M., Poulain, P.-M., and Zambianchi, E.: Validation of HF radar-derived currents in the Gulf of Naples with Lagrangian data, *IEEE J. Oceanic Eng.*, **41**, 569–586, 2016.

Kaplan, D. M., Largier, J., and Botsford, L. W.: HF radar observations of surface circulation off Bodega Bay (northern California, USA), *J. Geophys. Res.*, **110**, C10020, https://doi.org/10.1029/2005JC002959, 2005.

Kim, S. Y., Terrill, E. J., and Cornuelle, B. D.: Mapping surface currents from HF radar radial velocity measurements using optimal interpolation, *J. Geophys. Res.*, **113**, C10023, https://doi.org/10.1029/2007JC004244, 2008.

Lipa, B. J. and Barrick, D. E.: Least-squares methods for the extraction of surface currents from CODAR crossed-loop data: application at ARSLOE, *IEEE J. Oceanic Eng.*, **8**, 226–253, https://doi.org/10.1109/JOE.1983.1145578, 1983.

Lipa, B., Nyden, B., Ullman, D. S., and Terrill, E.: SeaSonde radial velocities: derivation and internal consistency, *IEEE J. Oceanic Eng.*, **31**, 850–861, 2006.

https://doi.org/10.5194/egusphere-2026-3013
Preprint. Discussion started: 5 June 2026

Lloyd, S.: Least squares quantization in PCM, *IEEE Trans. Inf. Theory*, **28**, 129–137, https://doi.org/10.1109/TIT.1982.1056489, 1982.

MacQueen, J.: Some methods for classification and analysis of multivariate observations, *Proc. 5th Berkeley Symp. Math. Stat. Probab.*, 1967.

Manso-Narvarte, I., Fredj, E., Jordà, G., Páez-Reyes, M., and Tintoré, J.: 3D reconstruction of ocean velocity from high-frequency radar and acoustic Doppler current profiler: a model-based assessment study, *Ocean Sci.*, **16**, 575–591, https://doi.org/10.5194/os-16-575-2020, 2020.

Paduan, J. D. and Cook, M. S.: Mapping surface currents in Monterey Bay with CODAR-type HF radar, *Oceanography*, **10**, 49–52, https://doi.org/10.5670/oceanog.1997.21, 1997.

Paduan, J. D. and Graber, H. C.: Introduction to high-frequency radar: reality and myth, *Oceanography*, **10**, 36–39, https://doi.org/10.5670/oceanog.1997.18, 1997.

Paduan, J. D., Kim, K. C., Cook, M. S., and Chavez, F. P.: Calibration and validation of direction-finding high-frequency radar ocean surface current observations, *IEEE J. Oceanic Eng.*, **31**, 862–875, https://doi.org/10.1109/JOE.2006.886195, 2006.

Paduan, J. D. and Washburn, L.: High-frequency radar observations of ocean surface currents, *Annu. Rev. Mar. Sci.*, **5**, 115–136, https://doi.org/10.1146/annurev-marine-121211-172315, 2013.

Roarty, H., Cook, T., Hazard, L., George, D., Harlan, J., Cosoli, S., Griffa, A., and Glenn, S.: The global high frequency radar network, *Front. Mar. Sci.*, **6**, 164, https://doi.org/10.3389/fmars.2019.00164, 2019.

Shi, J., Chen, X., Li, X., Chen, P., and Shao, Q.: Construction of sea surface current vectors using a single HF radar: a theoretical study, *iScience*, **25**, 105722, https://doi.org/10.1016/j.isci.2022.105722, 2022.

Stewart, R. H. and Joy, J. W.: HF radio measurements of surface currents, *Deep-Sea Res.*, **21**, 1039–1049, https://doi.org/10.1016/0011-7471(74)90066-7, 1974.

Tonani, M., Sykes, P., King, R. R., McConnell, N., Pequignet, A.-C., O'Dea, E., Graham, J. A., Polton, J., and Siddorn, J.: The impact of a new high-resolution ocean model on the Met Office North-West European Shelf forecasting system, *Ocean Sci.*, **15**, 1133–1158, https://doi.org/10.5194/os-15-1133-2019, 2019.

Wyatt, L.: High-frequency radar applications in coastal monitoring, planning and engineering, *Aust. J. Civ. Eng.*, **12**, 1–15, 2014.