remote sensing

MDPI

Article

Constraint-Based Adaptive Grids for Compact Representation of Marine Scalar Fields

Jinhong Park [1,*] and Sung Yong Kim [2]

[1] Department of Research and Development, Sfractum Co., Ltd., GA-529, 54, Changeop-ro, Sujeong-gu, Seongnam-si, Gyeonggi-do 13449, Republic of Korea
[2] Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, 291 Daehak-ro, Yuseong-gu, Daejeon-si 34141, Republic of Korea; syongkim@kaist.ac.kr
* Correspondence: jhpark@sfractum.com

Highlights

What are the main findings?

- A deterministic, scale-admissible dyadic representation selects its parameter profile from calibration fields and holds it fixed during the corresponding held-out tests, preventing evaluation-time retuning.
- Adaptive tile placement provided a distinct compactness–fidelity trade-off, while quadtree comparisons showed that similar aggregate losses can accompany different tile boundaries.

What are the implications of the main findings?

- Exact tile-to-cell membership makes the output an interpretable multiresolution layer for visualization, selective transfer, and exploratory analysis while leaving the source grid authoritative.
- Spatial organization and byte coding serve complementary purposes: direct scientific codecs favor payload reduction, whereas hybrid coding can reduce payload while retaining adaptive topology.

Abstract

Fixed-block representations can obscure heterogeneous local variability in marine scalar fields. We present a deterministic adaptive representation in which anchored dyadic tiles merge under a field-normalized local-range rule and a realized-tile diameter cap. Area-weighted tile values provide lossy reconstruction, whereas exact tile-to-cell membership preserves the link to the authoritative source grid. A profile selected within an evaluated candidate set by Pareto filtering and normalized minimax regret ($P_{\max} = 2$ and $S_{\text{th}} = 0.9737439$) is fixed before held-out evaluation, with partitions regenerated for each field. In East Sea temperature fields from the Korea Hydrographic and Oceanographic Agency (KHOA) Regional Oceanic Modeling System (ROMS), the method provided a compactness–fidelity compromise whose advantages relative to fixed-block references depended on the metric and block size. Quadtree comparisons showed that similar aggregate losses can accompany different tile boundaries, supporting a transparent representation and calibration framework rather than a universally superior split rule. Salinity and chlorophyll-a cases assessed broader applicability without implying universal parameter transfer. Direct compression with the ZFP and SZ3 scientific-data compressors achieved smaller serialized payloads under the stated benchmark protocol; hybrid coding reduced payload while retaining adaptive topology. The resulting derived layer is designed to

Check for updates

Academic Editor: Zhenwei Shi

Received: 14 July 2026
Revised: 8 September 2026
Accepted: 21 September 2026
Published: 1 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

support multiresolution visualization, selective transfer, and exploratory analysis without replacing the source grid.

Keywords: adaptive grid generation; local similarity; scale-admissible partitioning; derived spatial representation; quadtree comparison; KHOA ROMS; scientific data compression; ZFP; SZ3; marine environmental data

1. Introduction

1.1. Marine Environmental Fields as Derived Spatial Representations

Marine observing systems, numerical models, data-assimilation products, and Earth-observation services now generate gridded environmental fields at resolutions and update frequencies that impose substantial storage, transfer, rendering, and exploratory-analysis costs. The scientific interpretation of these fields depends on scale: patterns in remotely sensed or mapped variables change with spatial support, resolution, and aggregation rule [1–3]. Retaining every native cell may therefore be unnecessary for some visualization and delivery products, whereas uniform coarsening can obscure fronts, coastal transitions, plumes, and other locally heterogeneous features.

The distinction between an authoritative source grid and a derived representation is central to this study. The proposed adaptive grid (AG) is a post-processing representation of an existing scalar field. It neither changes the discretization used by a numerical model nor supplies the conservation, stability, and operator consistency required of a computational mesh. Instead, it stores an explicit set of spatial tiles, a representative value for each tile, and an exact membership map from tiles back to valid source cells. The representative values are lossy, whereas the partition membership can be stored losslessly. Intended uses include multiresolution visualization, selective spatial transfer, level-of-detail rendering, and exploratory queries; precise model diagnostics remain tied to the source grid.

This framing is relevant to marine digital twins, which must connect physical entities, data, services, and virtual models while addressing scalability, interoperability, and explainability [4–7]. Image-based ocean-feature systems likewise demonstrate the value of spatial representations that expose coherent structures for interactive inspection and temporal comparison [8]. Recent data-driven ocean-forecasting systems, including GLONET and FuXi-Ocean, further illustrate the growing scale and multivariate structure of global ocean products and the need for process-aware evaluation [9,10]. Those systems predict evolving ocean states; the present work instead constructs a downstream representation of an existing field. The primary experiment uses model-derived near-surface potential-temperature fields, and the external case study uses satellite-derived chlorophyll-a. The remote-sensing relevance lies in scale-aware representation of gridded marine information originating from Earth observation, models, or data fusion, rather than in sensor-level retrieval or onboard image coding.

1.2. Numerical Grids, Quadtrees, and Scientific Compression

Hierarchical partitioning has a long history. Quadtrees provide efficient two-dimensional subdivision and retrieval [11,12]; tree-traversal split-and-merge methods compact homogeneous image regions [13]; and quadtree decomposition has been used for image coding and remote-sensing segmentation [14,15]. Marine applications include variable-resolution bathymetric products and storage-saving multibeam maps [16,17]. Discrete global grid systems address a broader global indexing and geospatial-operation

problem [18]. These precedents establish the tree and indexing context, but the present contribution is not a new tree data structure.

Adaptive numerical meshes solve a different problem. Adaptive mesh refinement and variable-resolution ocean models allocate degrees of freedom while respecting numerical equations and coupling requirements [19–23]. The proposed post-processing tiles do not replace those grids. Similarly, recent remote-sensing encoding work uses content-adaptive segmentation, sampling allocation, and compressed-sensing reconstruction for image acquisition [24]; our method instead aggregates an existing physical scalar field under an explicit physical-scale cap and evaluates scientific field metrics.

Scientific visualization and compression research further shows that spatial resolution and value precision are complementary dimensions. A unified hierarchical layout can support progressive decoding along resolution–precision paths [25], and adaptive multilinear meshes jointly vary resolution and precision while retaining queryable structure [26]. Direct floating-point codecs pursue a different primary objective: the ZFP floating-point compressor transforms small blocks into compact fixed-rate or accuracy-controlled bit streams [27], whereas the SZ and SZ3 scientific-data compressors use modular prediction, quantization, and lossless coding under error controls [28,29]. Standardized benchmarking emphasizes consistent data layout, error metrics, and operating conditions [30], and quantity-of-interest preservation research cautions that pointwise error bounds alone may not protect downstream analyses [31]. Atmospheric-data studies similarly show that useful precision must be related to information content and end use [32].

Learned compression now spans Earth-observation and Earth-system-model data [33]. Recent remote-sensing studies have addressed fixed-quality neural coding and learned spectral–spatial transforms [34,35], while a convolutional-attention autoencoder has been developed for three-dimensional marine temperature fields [36]. Such methods can provide compact latent representations after training, but they do not necessarily expose a deterministic spatial partition or exact tile-to-cell membership. The present study therefore compares serialized payloads with ZFP and SZ3 while treating explicit topology and dense numerical value compression as distinct, potentially complementary capabilities.

1.3. Study Contributions and Scope

This study makes four contributions. First, it defines a deterministic, constraint-based adaptive representation of a curvilinear masked field: globally anchored dyadic tiles are merged by a field-normalized local-range rule, represented by values weighted using center-derived polygon areas, and linked to the source grid through exact tile-to-cell membership. Second, it implements physical-scale admissibility as a block-level safeguard based on the maximum World Geodetic System 1984 (WGS84) geodesic distance among the six corner pairs of every candidate tile. Third, it separates calibration from evaluation: one parameter profile is selected from four calibration fields by four-objective Pareto filtering and normalized minimax regret, then held fixed while a field-specific partition is regenerated for every within-sequence held-out field. Fourth, it distinguishes spatial representation from numerical value coding through uniform and quadtree comparators, active-cap sensitivity, cross-variable and cross-region case studies, measured serialized bytes, direct scientific codecs, hybrids, and harmonized timing boundaries.

The method is a derived spatial layer rather than a numerical mesh, and representative values provide a lossy reconstruction while tile membership remains exact. Its empirical scope comprises twelve nominal monthly first-day East Sea fields, same-grid salinity, and a satellite-derived chlorophyll-a case study. These tests examine within-sequence transfer and two forms of cross-dataset application without implying unrestricted geographic, variable, or sensor transfer. The physical cap remains part of the method as a safeguard,

although it is nonbinding throughout the evaluated East Sea temperature candidate set. The unconstrained-range quadtree is treated as a separately calibrated comparator because both its maximum level and threshold differ; it does not isolate the effect of removing a single component.

The remainder of the paper follows the operational sequence. Section 2 fixes the datasets, calibration–held-out roles, exact partition rule, objectives, comparators, and measurement boundaries before results are examined. Section 3 moves from parameter response and spatial behavior to within-sequence held-out evaluation, realized-cap sensitivity, cross-dataset validation, monthly first-day topology variation, and fixed-budget model-mask-boundary allocation. Section 4 interprets the representation and then evaluates byte-level capability relative to direct codecs and hybrids. Reproducibility, physical metrics, alternative codec decisions, runtime protocols, and hybrid operating spaces are documented in the Appendices; further diagnostics and unvalidated extension formulas are placed in the Supplementary Materials.

2. Materials and Methods

2.1. Study Data and Calibration–Held-Out Design

The primary source comprises twelve two-dimensional, near-surface (1 m) potential-temperature fields (hereafter, temperature) from the public Korea Hydrographic and Oceanographic Agency (KHOA) Regional Oceanic Modeling System (ROMS) numerical-prediction archive [37,38]. The source `temp` variable is reported in degrees Celsius and described as time averaged. We label the ordered fields M01–M12 and use their filename day indices to identify nominal monthly first-day samples within an annual cycle. The retained files do not resolve the calendar year, averaging interval, or forecast lead, so no finer temporal interpretation is made. We intentionally use these twelve fields for a concise methodological comparison; daily and continuous time-series behavior remains untested. Each array has 322×706 cells, and the common model-ocean mask contains $K_0 = 130{,}131$ finite valid cells spanning approximately 126.9209–150.1657°E and 31.5231–51.9262°N. Invalid or land cells are excluded from partitioning and fidelity metrics.

The same files also contain a near-surface `salt` variable on the identical grid and mask. This same-grid variable supplies two distinct tests: direct application of the temperature-selected profile without recalibration and independent calibration of a salinity profile. The distributed East Sea inputs are losslessly compressed variable subsets retaining `temp`, `salt`, `lon`, `lat`, and `mask` at their original float64 precision and array dimensions. No spatial resampling or additional time or depth extraction is performed. The 1 m designation comes from the supplied source filenames; the files do not contain an independent depth coordinate.

For an external cross-region, cross-grid, and observation-source case study, we use the National Oceanic and Atmospheric Administration (NOAA) Suomi National Polar-orbiting Partnership (S-NPP) Visible Infrared Imaging Radiometer Suite (VIIRS) science-quality Level-3 monthly chlorophyll-a product [39,40]. Twelve 2024 monthly composites were downloaded for the California Current region (30–45°N, 135–115°W) at approximately 4.17 km resolution, yielding a 401×535 array per month. For positive concentration C, the primary field is the dimensionless quantity $\log_{10}[C/(1\ \mathrm{mg\,m^{-3}})]$. Its common calibration support is the intersection of valid pixels in January, April, July, and October; each held-out month is evaluated on that support intersected with its own valid pixels. This procedure prevents held-out missingness from defining the primary calibration support. An all-12-month common-support analysis and an untransformed concentration analysis are reported as sensitivities.

For every dataset, M01, M04, M07, and M10 were designated as calibration fields before candidate evaluation. M02, M03, M05, M06, M08, M09, M11, and M12 were reserved for within-sequence held-out evaluation. Candidate bounds, Pareto membership, normalization extrema, tie breaking, and all calibration-fixed comparator settings were determined without held-out outcomes. Three explicitly field-specific rate comparators use each evaluation field's proposed tile count: the nearest tile-count uniform baseline chooses among discrete block sizes, whereas the field-retuned and $P_{\max}$-matched local root mean square error (RMSE) variants adjust their tolerances. These are descriptive rate controls and are not evidence of parameter transfer. The selected parameter profile is fixed, but the spatial partition is not: the deterministic algorithm regenerates a partition from each evaluation field. The eight held-out observations are correlated positions within one sequence and are not treated as statistically independent.

Table 1 summarizes the data sources, transformations, calibration–held-out assignments, and role of each evaluation branch.

Table 1. Datasets, sampling design, and validation roles. The East Sea sequence contains twelve nominal monthly first-day fields; the VIIRS fields are dated monthly composites.

Dataset/Source	Variable and Transform	Grid/Support	Calibration–Held-Out Split	Role and Scope
KHOA ROMS, East Sea	Near-surface (1 m) potential temperature; Celsius; no transform	322 × 706 curvilinear grid; 130,131 common valid model-ocean cells; 126.9209–150.1657°E, 31.5231–51.9262°N	M01/M04/M07/M10 calibration; remaining eight nominal monthly first-day fields held out	Primary within-sequence evaluation and codec study
KHOA ROMS, East Sea	Near-surface `salt`; no transform	Same grid and model mask as temperature	Same four-to-eight split	Direct temperature-profile transfer and salinity-specific calibration
NOAA S-NPP VIIRS, California Current, 2024	Monthly chlorophyll-a; primary $\log_{10}[C/(1\ \mathrm{mg\,m^{-3}})]$	401 × 535 regular grid; calibration- valid intersection followed by field-valid intersection	Jan/Apr/Jul/Oct calibration; remaining eight monthly composites held out	Cross-region, cross-grid, satellite-source case study; raw-field and common-mask sensitivities
Proposed adaptive-grid profile-transfer analyses	K/K_0, RMSE_w/R, D_{vario}, and $1 - G_{\mathrm{grad}}$ for cross-variable comparisons	Complete nonoverlap, finite representatives, exact membership, and scale admissibility	Profiles fixed within each transfer test; partitions regenerated field by field	Held-out summaries describe the sampled sequences rather than an independent population

2.2. Derived Representation and Operational Workflow

For a scalar field on an $n_r \times n_c$ grid, let $\Omega = \{(r, c) : 1 \le r \le n_r,\ 1 \le c \le n_c\}$ and let $m_{rc} \in \{0, 1\}$ denote the applicable validity mask. The valid domain is

$$\Omega_v = \{(r, c) \in \Omega : m_{rc} = 1,\ z_{rc} \text{ is finite}\}. \tag{1}$$

An adaptive representation for field z is a partition $\mathcal{T} = \{T_k\}_{k=1}^{K}$ of Ω_v, together with a representative value $\bar{z}_k$ and tile metadata sufficient to reproduce the exact source-cell membership of every tile. Reconstruction assigns $\widehat{z}_{rc} = \bar{z}_k$ for $(r, c) \in T_k$. Thus $\mathcal{T}$ is exact as an index relation, but $\widehat{z}$ is generally a lossy piecewise-constant field. This dataset-independent definition covers both the 322 × 706 curvilinear KHOA grid and the regular VIIRS subset.

Figure 1 separates offline calibration, fixed-profile application, spatial and cross-dataset validation, and scientific-codec evaluation. Adaptive selection uses the structural tile ratio and three field-fidelity objectives. Serialized bytes are evaluated afterward and do not enter the four-objective adaptive-grid minimax rule.

2.3. Physical Coordinates and Grid Geometry

Longitude and latitude are not treated as Euclidean coordinates. For each common grid, valid-cell centers are transformed from WGS84 longitude–latitude to an azimuthal equidistant (AEQD) projection centered on the valid domain. The projected coordinates (x_i, y_i), in metres, provide local approximations for distances and physical gradients. Adjacent center-to-center spacings are also computed geodesically on WGS84 along the two curvilinear index directions. With their valid medians denoted by d_ξ and d_η, the conservative representative spacing is

$$d = \max(d_\xi, d_\eta). \tag{2}$$

For M01, $d_\xi = 3214.894$ m, $d_\eta = 3342.469$ m, and $d = 3342.469$ m.

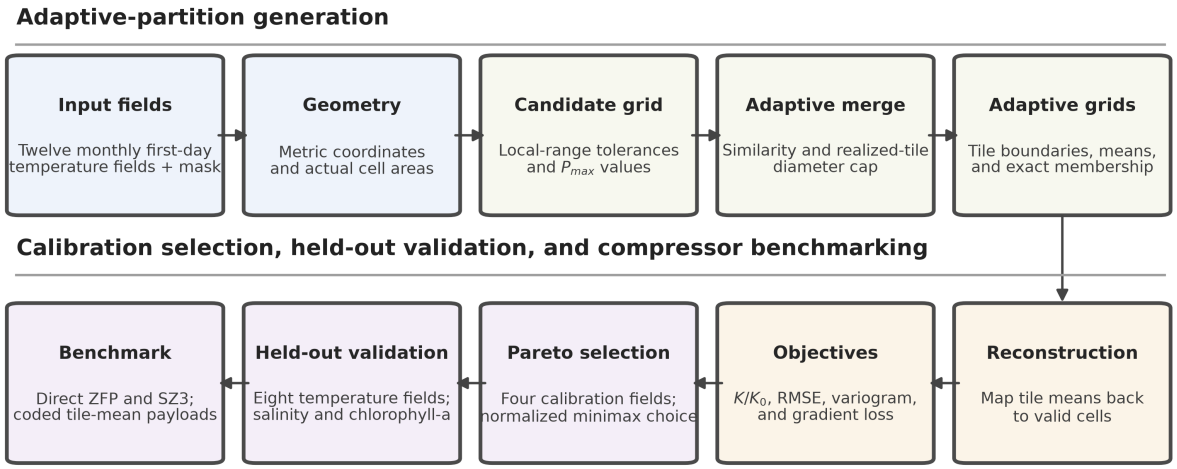


Figure 1. Analysis workflow and pipeline. Four calibration fields determine one adaptive parameter profile by exact realized-tile scale screening, Pareto filtering, and normalized minimax regret. The profile is fixed before field-specific partitions are regenerated for eight within-sequence held-out fields. Same-grid salinity, external satellite chlorophyll-a, spatial comparators, direct codecs, hybrids, and runtime protocols are evaluated in distinct branches. The output is a lossy derived field with an explicit partition and exact tile-to-cell membership rather than a replacement numerical-model grid.

Area weights are not approximated by a cosine-latitude factor. For each longitude and latitude array, every interior cell corner is estimated as the arithmetic mean of its four surrounding cell centers. Boundary corners are obtained by linear extrapolation from the adjacent center-pair midpoints, and the four terminal corners are continued from their neighboring corner values. The resulting curvilinear corner polygons are represented in a Lambert azimuthal equal-area (LAEA) projection centered on the domain, and their polygon areas define $A_i > 0$. For M01 the valid-cell areas range from 7.556×10^6 to 1.417×10^7 m^2, with median 1.074×10^7 m^2. AEQD is used where local metric distance or derivatives are required; LAEA is used where area preservation is required.

2.4. Scale-Admissible Adaptive Partition Construction

The parameter pair is defined consistently as $\theta = (P_{\max}, S_{\text{th}})$. For field m, the global robust amplitude range and derived aggregation tolerance are

$$R_m = Q_{97.5}(z_m) - Q_{2.5}(z_m), \qquad \tau_m = (1 - S_{\text{th}})R_m, \tag{3}$$

where quantiles are computed over Ω_v. A small numerical floor is applied to R_m to handle constant fields. The fixed profile specifies a dimensionless tolerance: S_{th} remains unchanged during held-out application, whereas R_m and hence the absolute tolerance τ_m are computed

from each input field. The robust global normalization limits the influence of extreme field values on the scale of τ_m; it does not make the blockwise maximum–minimum statistic itself robust.

Dyadic level p has side length $s_p = 2^p$ cells. Blocks are aligned to the global top-left array origin and traversed deterministically from coarse to fine, $p = P_{\max}, \ldots, 1$. A block is eligible only if it is complete, all of its cells are valid, and none has previously been assigned. Because the grid has only 322 rows, a $2^9 \times 2^9 = 512 \times 512$ complete block cannot occur; $P_{\max} = 9$ would therefore duplicate the $P_{\max} = 8$ traversal under the strict complete-block policy and was excluded from the search. An eligible block is accepted when

$$\Delta_m(T) = \max_{i \in T} z_{m,i} - \min_{i \in T} z_{m,i} \leq \tau_m. \tag{4}$$

Incomplete blocks at the right or bottom array boundary, blocks crossing the mask, and residual unassigned valid cells are refined; any remaining valid cells become singleton tiles. The procedure guarantees complete nonoverlapping coverage of Ω_v, but it does not impose 2:1 balance and is not asserted to solve a globally optimal partition problem.

The representative tile value is the mean weighted by the center-derived polygon areas

$$\bar{z}_{m,k} = \frac{\sum_{i \in T_k} A_i z_{m,i}}{\sum_{i \in T_k} A_i}. \tag{5}$$

The implementation scans rectangular grid arrays at each level. Its worst-case work is $O(n_r n_c P_{\max})$ and its memory requirement is $O(n_r n_c)$. For bounded $P_{\max}$, the work is linear in the total array size; expressing this in terms of valid-cell count additionally assumes a valid fraction bounded away from zero.

To constrain representational scale, a practical correlation range $\hat{r}_m$ is estimated in physical coordinates as described below. The calibration-only cap is fixed as

$$D_{\text{cap}} = \alpha \min_{m \in \mathcal{C}} \hat{r}_m, \qquad \alpha = 0.7, \tag{6}$$

where $\mathcal{C}$ is the four-field calibration set. For candidate block T, let $\boldsymbol{q}_{T,1}, \ldots, \boldsymbol{q}_{T,4}$ be its four outer cell-corner coordinates obtained from the curvilinear grid, and let d_{WGS84} denote ellipsoidal geodesic distance. Its exact implemented diameter is

$$D(T) = \max_{1 \leq a < b \leq 4} d_{\text{WGS84}}(\boldsymbol{q}_{T,a}, \boldsymbol{q}_{T,b}). \tag{7}$$

A merge is admissible only if $D(T) \leq D_{\text{cap}}$; otherwise traversal continues to finer levels. Thus the physical rule is evaluated block by block and does not substitute 2^p times one representative spacing for the realized geometry. The value $\alpha = 0.7$ was fixed before held-out evaluation as a prespecified design margin that places the cap 30% below the shortest estimated calibration practical range. It is a design choice rather than an optimized constant, confidence bound, or universal physical coefficient. The search envelope $P_{\max} \in \{0, \ldots, 8\}$ is likewise part of the study design. Cap-activity diagnostics are reported because the selected East Sea temperature profile is nonbinding at its correlation-derived cap.

2.5. Fidelity Objectives and Calibration-Only Selection

2.5.1. Area-Weighted Reconstruction Error

The first field loss is the area-weighted root mean square error

$$\text{RMSE}_{w,m} = \left[\frac{\sum_{i \in \Omega_v} A_i (z_{m,i} - \hat{z}_{m,i})^2}{\sum_{i \in \Omega_v} A_i} \right]^{1/2}. \tag{8}$$

2.5.2. Normalized Semivariogram Discrepancy

Spatial-correlation fidelity is evaluated after fitting a least-squares plane to the original sampled field in AEQD coordinates. The same fitted plane is subtracted from the original and reconstructed samples, so reconstruction-specific trend fitting cannot conceal representation error. A deterministic common sample of up to 500 valid cells is drawn by selecting one seeded valid cell per array stratum, with the numbers of row and column strata determined from the array aspect ratio and target sample count, followed by seeded completion when required (seed 42). This procedure yields 15×34 strata for the East Sea grid. The lag cutoff $h_{\max}$ is half the diagonal of the axis-aligned bounding box of the sampled cell centers in AEQD coordinates. Common pairs with positive separation below $h_{\max}$ are assigned to 16 equal-width metric-distance bins. Empirical estimates are computed for every populated bin, while at least 30 pairs are required for a bin to enter the exponential-model fit. For differences $\delta_{ab} = e(\boldsymbol{s}_a) - e(\boldsymbol{s}_b)$ and N_h pairs in bin h, the Cressie–Hawkins estimator [41] is

$$\widehat{\gamma}(h) = \frac{1}{2} \frac{\left[N_h^{-1} \sum_{(a,b)\in h} |\delta_{ab}|^{1/2} \right]^4}{0.457 + 0.494/N_h + 0.045/N_h^2}. \tag{9}$$

This estimator is distinct from the subsequent weighted model fit [42–44]. The empirical semivariogram is normalized by the variance of the original detrended residual, $\widetilde{\gamma} = \widehat{\gamma}/\sigma_e^2$. An exponential model

$$\gamma_{\exp}(h) = c_0 + c\{1 - \exp(-3h/r)\} \tag{10}$$

is fitted by pair-count-weighted bounded nonlinear least squares. With semivariance normalized as above, the parameter bounds are $0 \le c_0 \le 3$, $0 \le c \le 5$, and $h_{\max}/100 \le r \le 2h_{\max}$, using the lag cutoff $h_{\max}$ defined above; optimization is limited to 30,000 function evaluations. A fit is flagged as boundary limited when $r \le 1.02(h_{\max}/100)$ or $r \ge 1.98h_{\max}$. Covariance-based standard errors are recorded when available, but they are not propagated as confidence intervals for the scale cap. The fitted r supplies the practical range used to define Equation (6). Directional fits are retained only as anisotropy diagnostics; candidate selection uses the omnidirectional range. The projected-coordinate, realized-tile, cap, and M01 variogram diagnostics are reported in Appendix B.

Using the bins populated for both the original and reconstructed residuals and n_h equal to the smaller available pair count, the dimensionless discrepancy is

$$D_{\text{vario},m} = \left[\frac{\sum_h n_h \{\widetilde{\gamma}_{\widehat{z}_m}(h) - \widetilde{\gamma}_{z_m}(h)\}^2}{\sum_h n_h} \right]^{1/2}. \tag{11}$$

2.5.3. Physical-Gradient Magnitude Fidelity

Index-direction derivatives of z, x, and y are computed by centered differences when both neighbors are valid and by a one-sided difference when only one neighbor is valid. No stencil crosses the invalid mask. The evaluation domain $\Omega_g \subseteq \Omega_v$ contains cells with both index-direction derivatives, a finite nonsingular coordinate Jacobian, and valid original and reconstructed gradients. Solving the Jacobian yields projected physical-gradient components in units of the variable per metre. Let $g_i = \|\nabla z_i\|$ and $\widehat{g}_i = \|\nabla \widehat{z}_i\|$. The magnitude-sensitive similarity is

$$G_{\text{grad},m} = \text{clip}_{[0,1]} \left(1 - \frac{\sum_{i\in\Omega_g} A_i |g_i - \widehat{g}_i|}{\sum_{i\in\Omega_g} A_i (g_i + \widehat{g}_i)} \right), \tag{12}$$

with similarity defined as one when the denominator is zero over Ω_g. The optimization loss is $1 - G_{\mathrm{grad},m}$. Unlike a correlation coefficient, Equation (12) penalizes systematic attenuation or amplification of gradient magnitude.

2.5.4. Pareto Filtering and Minimax Compromise

Candidate enumeration was fixed from the calibration fields alone. The broad search used every $P_{\max} \in \{0, \ldots, 8\}$ and the threshold set

$$\mathcal{S}_{\mathrm{broad}} = \{0, 0.25, 0.50, 0.75, 0.90, 0.95, 0.97, 0.975, 0.98, \\ 0.985, 0.99, 0.9925, 0.995, 0.9975, 0.999, 0.9995, 0.9999, 1\}.$$

Its provisional minimax solution defined the neighboring bracket $[0.97, 0.98]$, which was sampled at 17 equally spaced thresholds. Exact normalized block-range transitions were then pooled from M01, M04, M07, and M10 over the one-level neighborhood of the provisional $P_{\max}$. Twenty-five quantiles of that pool and the 10 transitions nearest the provisional threshold in each calibration field were retained, values were rounded to seven decimal places, and duplicates were removed. The union comprised 94 thresholds and nine maximum levels, or 846 parameter pairs and 3384 field-level calibration evaluations. Table S1 summarizes the successive evaluated designs, and the accompanying candidate files retain every evaluated coordinate.

The three empirical fidelity quantities are objectives, not hard feasibility thresholds. For the calibration set $\mathcal{C} = \{\mathrm{M01}, \mathrm{M04}, \mathrm{M07}, \mathrm{M10}\}$, four worst-observed objectives are minimized:

$$F_1(\theta) = \max_{m \in \mathcal{C}} K_m / K_{0,m}, \qquad F_2(\theta) = \max_{m \in \mathcal{C}} \mathrm{RMSE}_{w,m}, \\ F_3(\theta) = \max_{m \in \mathcal{C}} D_{\mathrm{vario},m}, \qquad F_4(\theta) = \max_{m \in \mathcal{C}} (1 - G_{\mathrm{grad},m}). \tag{13}$$

The maxima are observed calibration summaries, not universal error bounds. After the hard structural and scale rules, dominated candidates are removed. Let $\mathcal{P}$ denote the resulting four-objective Pareto set. Each objective is normalized only within $\mathcal{P}$:

$$\widetilde{F}_j(\theta) = \begin{cases} \dfrac{F_j(\theta) - F_j^{\min}}{F_j^{\max} - F_j^{\min}}, & F_j^{\max} > F_j^{\min}, \\ 0, & F_j^{\max} = F_j^{\min}. \end{cases} \tag{14}$$

The selected profile minimizes its worst normalized regret:

$$\theta^* = \arg\min_{\theta \in \mathcal{P}} \max_{j \in \{1,\ldots,4\}} \widetilde{F}_j(\theta). \tag{15}$$

Ties are resolved in order by minimum mean normalized regret, minimum tile ratio, smaller $P_{\max}$, and larger S_{th}. This is a prespecified compromise within the evaluated candidate set, not a continuous-parameter optimum or an application-independent utility function. The staged threshold set samples partition transitions rather than exhaustively enumerating them.

2.6. Spatial Comparators and Fair Matching

All primary spatial methods share the ocean mask, global dyadic anchoring, strict edge policy, polygon-area-weighted representatives, reconstruction rule, and metric implementation. They differ in split criterion and in when parameters are adjusted. Table 2 makes these distinctions explicit. The Field-retuned local-RMSE quadtree answers a rate-matched

descriptive question by adjusting its tolerance separately for every field and therefore does not test parameter transfer. The calibration-fixed comparator addresses transfer directly. The unconstrained-range quadtree uses the same local-range merge statistic at the prespecified structural ceiling $P_{\max} = 8$, with a separately calibrated threshold and no realized-diameter safeguard. Because more than one component changes, it is treated as a comparator rather than a component ablation. Both calibration-fixed quadtree families use the same four-objective Pareto–minimax selection convention within their own candidate sets; their local split criteria, rather than their calibration objectives, distinguish them from the proposed construction.

The nearest tile-count uniform baseline accepts fully valid blocks at one fixed dyadic size and leaves valid cells in mask-crossing or incomplete blocks as singletons. It therefore combines fixed-size interior blocks with singleton boundary cells; it does not assign one physical scale everywhere. The available tile counts need not decrease monotonically with block size. Fixed 2×2 and 4×4 references are additionally reported to show how conclusions depend on this discrete nearest-count choice.

Table 2. Spatial comparators and matching rules.

Method	Construction	Parameter Timing	Question Answered
Nearest tile-count uniform baseline	Fully valid blocks at one dyadic size plus singleton residual cells; nearest available tile count	Per field from discrete sizes	What trade-off does this strict-mask nearest-count rule produce?
Field-retuned local-RMSE quadtree	Area-weighted local-RMSE split tolerance binary-searched to the proposed tile count	Retuned on every field	At nearly equal tile count, what fidelity does an RMSE-oriented quadtree achieve?
Calibration-fixed local-RMSE quadtree	Local-RMSE splitting; one normalized tolerance selected by calibration Pareto–minimax	Fixed on held-out fields	How does a profile using local-RMSE splitting transfer?
$P_{\max}$-matched local-RMSE quadtree	Local-RMSE splitting with the selected proposed $P_{\max}$ and per-field tile matching	$P_{\max}$ fixed; tolerance per field	How much difference remains after sharing the maximum dyadic level?
Unconstrained-range quadtree	Same local block-range rule with an unconstrained structural maximum level and one calibration-selected threshold	Distinct maximum level and threshold fixed before evaluation; no realized-diameter filter	What behavior is obtained from a separately calibrated range-rule quadtree without the physical safeguard?
Proposed adaptive grid	Range-based dyadic partition, exact realized-diameter safeguard, and four-objective calibration minimax	Parameter profile fixed before held-out evaluation; partition regenerated per field	Can an explicit balanced operating profile transfer across the tested fields?

An operating set is the collection of distinct partitions realized while one shared calibration parameter is varied. The local-RMSE structural-ceiling and proposed-depth sweeps each vary one normalized tolerance shared across M01, M04, M07, and M10. The structural-ceiling sweep fixes $P_{\max} = 8$, the prespecified structural ceiling on the 322-row grid, and therefore asks what the local-RMSE rule can attain without the selected depth restriction. The proposed-depth sweep instead fixes $P_{\max} = 2$, the depth selected for the proposed profile, and isolates the effect of evaluating the local-RMSE rule at that same depth. These shared-threshold operating sets are distinct from the Field-retuned and proposed-depth rate-matched policies: the latter adjust their tolerances separately for each field to approach the proposed tile count and therefore appear only as evaluated points. The proposed-depth rate-matched policy is the same $P_{\max}$-matched local-RMSE comparator

listed in Table 2; the descriptive figure label omits the parameter symbol for legibility. Likewise, the unconstrained-range label in the figure denotes the Unconstrained-range comparator in Table 2. Any normalized regret reported for a candidate family is defined within that family and is not compared across algorithms.

2.7. Cross-Variable and External-Observation Validation Methods

The salinity and chlorophyll-a experiments distinguish direct temperature-profile transfer without recalibration from target-specific calibration. In the former, the complete East Sea temperature profile $(P_{\max}, S_{\text{th}}, D_{\text{cap}}) = (2, 0.9737439, 366.9376\ \text{km})$ is applied directly, and a new partition is generated from every target-variable field. In the latter, the same algorithm, four objectives, calibration months, and held-out policy are retained, but the cap is derived and $(P_{\max}, S_{\text{th}})$ is selected using only the four calibration fields of the target dataset. The target-specific profiles were $(1, 0.380145, 1084.3696\ \text{km})$ for salinity and $(1, 0.8875, 1744.5918\ \text{km})$ for primary $\log_{10}[C/(1\ \text{mg}\,\text{m}^{-3})]$ chlorophyll-a. The calibration-fixed local-RMSE and unconstrained-range thresholds use the corresponding target dataset's calibration fields. The uniform baseline instead selects a discrete block size separately for each field by its nearest tile count to the target-specific proposed representation.

Because raw RMSE has variable-dependent units and amplitudes, the cross-dataset table reports RMSE_w/R, where R is the robust field range from Equation (3), together with K/K_0, D_{vario}, and $1 - G_{\text{grad}}$. Raw RMSE_w and serialized-record quantities remain in the detailed Supplementary Data. For VIIRS, the primary valid support is fixed from calibration months before held-out evaluation, then intersected with each held-out field's valid pixels; this preserves observed missingness without allowing a held-out month to determine the calibration support.

The chlorophyll-a variogram fits require a specific qualification. All four calibration fits reached the permitted range upper bound of $2h_{\max}$, twice the lag cutoff. Their practical ranges are therefore unresolved within the sampled domain. Chlorophyll-a is used to test deterministic representation across a new region, regular grid, variable, and satellite source, but not to claim empirical support for correlation-range scale selection. Cap activity is reported separately from predictive or reconstruction performance in every dataset.

2.8. Monthly First-Day Topology Diagnostics

For a tile with top-left array index (r_0, c_0), height h, and width w, the exact rectangle signature is (r_0, c_0, h, w). The signature assigned to a cell is that of its containing tile. For each of the 11 adjacent pairs among M01–M12, the cell partition-change fraction is the fraction of cells on their common valid support whose signatures differ. If $\mathcal{R}_m$ denotes the set of unique tile-rectangle signatures in field m, exact tile-rectangle similarity is the Jaccard index $|\mathcal{R}_m \cap \mathcal{R}_{m+1}|/|\mathcal{R}_m \cup \mathcal{R}_{m+1}|$. Cellwise modal-level persistence is the largest count attained by any dyadic tile level across the twelve fields, divided by 12; a tie between modal levels does not change this persistence value. These diagnostics describe the ordered monthly first-day fields and do not resolve intervening daily variation.

2.9. Fixed-Budget Model-Mask-Boundary Allocation Methods

We use M01 to test whether the same partition rule can express a prescribed spatial priority without changing the total tile-count budget. This calibration-field mechanism experiment is separate from profile selection and does not treat a geographic coastline as a fidelity target. Let $\chi_i \in \{0, 1\}$ denote the source-array validity mask. The in-domain valid–invalid interface is

$$\mathcal{B} = \{i \in \Omega_v : \exists j \in \mathcal{N}_8(i) \cap \Omega \text{ with } \chi_j = 0\}, \tag{16}$$

where neighbors outside the source array are ignored. For valid-cell center $\boldsymbol{s}_i$, the distance field is

$$d_i = 10^{-3} \min_{b \in \mathcal{B}} \left\| \pi_{\mathrm{AEQD}}(\boldsymbol{s}_i) - \pi_{\mathrm{AEQD}}(\boldsymbol{s}_b) \right\|_2, \tag{17}$$

where the AEQD coordinates are in metres and the factor 10^{-3} expresses d_i in kilometres. Natural Earth linework is used only as geographic context in the figure.

The predefined strength set is $\rho \in \{0, 0.20, 0.40, 0.60\}$. It spans no weighting through a 60% interface-local tightening because $w_\rho(0) = 1 - \rho$ remains 0.40 in the strongest case. With $\ell = 21.29$ km, the maximum WGS84 diameter among full-valid M01 blocks at the selected $P_{\max} = 2$, the priority relaxes over one coarsest admissible full-block length. The cellwise priority factor is

$$w_\rho(d) = 1 - \rho \exp(-d/\ell). \tag{18}$$

For candidate tile T, the conservative block tolerance is

$$\tau_T(\rho, c_\rho) = c_\rho \tau_{\mathrm{M01}} \min_{i \in T} w_\rho(d_i), \qquad \Delta(T) \le \tau_T(\rho, c_\rho), \quad D(T) \le D_{\mathrm{cap}}. \tag{19}$$

The global reference is $\rho = 0$ and $c_0 = 1$. For each positive ρ, deterministic bisection over c_ρ minimizes the absolute difference from the global-reference count $K_{\mathrm{ref}} = 27{,}513$; ties select the smaller sampled multiplier. All three tested strengths attain the target count exactly. A $\rho = 0.60$, $c_\rho = 1$ same-base-tolerance case is retained only as a cost diagnostic because it is allowed to use more tiles.

Distance strata are four equal-valid-ocean-area quartiles computed with LAEA cell areas. A tile that crosses a stratum boundary is prorated through the area-equivalent count

$$K_R^{\mathrm{eq}} = \sum_{T \in \mathcal{T}} \frac{A(T \cap R)}{A(T)} = \sum_{i \in R} \frac{A_i}{A_{T(i)}}, \qquad a_R = \frac{K_R^{\mathrm{eq}}}{K}. \tag{20}$$

Local RMSE_w is area weighted within each stratum. Original and reconstructed physical-gradient magnitudes are first computed on the full valid mask and then aggregated by stratum; D_{vario} remains a global diagnostic because a quartile-local semivariogram would introduce a different sampling design.

2.10. Evaluation, Byte Accounting, and Timing Protocols

Structural compactness is reported as K/K_0 and tile reduction as $1 - K/K_0$. Neither is treated as a byte-compression ratio. For dataset m and method or operating point h, let $S_{m,h}$ be the measured serialized or compressed payload size, $K_{0,m} = |\Omega_{v,m}|$ the number of valid source cells, and $b_m = 4$ bytes the size of one canonical float32 benchmark value. The serialized byte ratio is

$$B_{m,h} = \frac{S_{m,h}}{K_{0,m} b_m}. \tag{21}$$

For the East Sea data, the denominator is $130{,}131 \times 4 = 520{,}524$ bytes. The source Network Common Data Form (NetCDF) scalars are float64, and analysis uses float64 arithmetic; float32 is the common value representation used for serialization and codec input. Codec-comparison fidelity is measured after decoding float32 values, so it can differ slightly from spatial-construction metrics computed from float64 tile means. Thus B compares payloads with a canonical valid-cell float32 array, not with the original NetCDF file size. Shared coordinates and mask are excluded from payload accounting. Adaptive and hybrid decoding also requires the shared array shape and mask to expand tile records, whereas the direct-codec wrapper retains its array shape. Each payload includes its own measured wrapper bytes. Adaptive payloads include the tile-record information required

by their schema and representative values; hybrids preserve the same geometry fields and membership scope losslessly and compress only the representative-value stream.

Direct ZFP 1.0.1 candidates use fixed-accuracy and fixed-rate modes [27]. SZ3/pysz 1.0.3 candidates use absolute-error modes with INTERP, `INTERP_LORENZO`, and `LORENZO_REG` configurations [28,29]. Because these codecs require dense arrays, invalid cells are filled by nearest-valid-ocean values in array-index distance before coding, while all fidelity metrics remain restricted to the original ocean mask. This fill rule is part of the benchmark protocol and can affect both payload and reconstruction near the mask. Following fair-benchmark principles [30], the same fields, mask, denominator, and metric implementation are used. The primary direct-codec point minimizes the maximum calibration byte ratio among candidates no worse than the selected adaptive profile on all three calibration fidelity maxima. This is a matching rule conditional on the stated dense-fill protocol, not a universal tolerance or a guarantee for held-out fields. Primary AG+ZFP and AG+SZ3 hybrids use an absolute representative-value tolerance of 0.005 and are described as near-fidelity points rather than exact matches. The value 0.005 was fixed for the held-out analysis as an illustrative near-fidelity setting; no held-out loss was used to select it. Alternative common-byte-budget and method-specific minimax rules are documented in Appendix C.

Held-out point estimates are medians across eight fields. Relative method changes are computed for each paired field before taking their median; they are not ratios of separately aggregated medians. Uncertainty bars are percentile intervals from 5000 bootstrap resamples of the eight field identities using recorded deterministic seeds. They describe sensitivity to this limited set of fields and should not be interpreted as population-level confidence under independent temporal sampling. When the same proposed-profile fidelity statistic appears in both the spatial and codec comparisons, the spatial-analysis interval is reused rather than independently re-estimated with a second Monte Carlo draw.

Three runtime boundaries are kept separate. One-shot encoding starts from an in-memory source field plus shared mask/geometry and ends at an in-memory payload. Prepared-representation encoding starts from a prepared dense array or fixed adaptive partition. Common-output decoding starts from the payload plus required shared mask and shape and ends at a dense float32 ocean reconstruction. Each field–method–phase combination uses five warm-ups and 50 timed repeats, randomized method order, and disabled garbage collection. NetCDF input/output (I/O), calibration search, and fidelity metrics are excluded. Because the direct codecs are compiled while adaptive metadata and reconstruction use Python paths, the timings are implementation diagnostics rather than language-neutral speed rankings.

3. Results

3.1. Parameter Response and Fixed-Profile Selection

The final union search contained 94 distinct values of S_{th} and nine values of $P_{\max}$, giving 846 parameter pairs and 3384 calibration field evaluations. The search included both similarity endpoints and progressive refinements around changes in realized partition. After structural and physical-scale screening, 193 candidates were nondominated in the four calibration-max objectives. The selected candidate was $P_{\max} = 2$ and $S_{\mathrm{th}} = 0.9737439$. Its maximum normalized regret was 0.1957 and its mean normalized regret was 0.1235. The calibration maxima at this threshold were $K/K_0 = 0.2231$, $\mathrm{RMSE}_w = 0.0835$, $D_{\mathrm{vario}} = 0.0078$, and $1 - G_{\mathrm{grad}} = 0.1424$.

Figure 2 shows why neither a single metric nor the most compact candidate defines the chosen profile. Increasing the tolerance generally reduces tile count while increasing one or more fidelity losses, but discrete changes in the realized partition create plateaus and transitions. At the selected structural level limit, the normalized responses do not move

identically. The selected point limits its worst normalized regret rather than minimizing any one raw objective. Alternative calibration-only decision rules move to visibly different spatial-rate regimes, which motivates reporting the rule as part of the model definition.

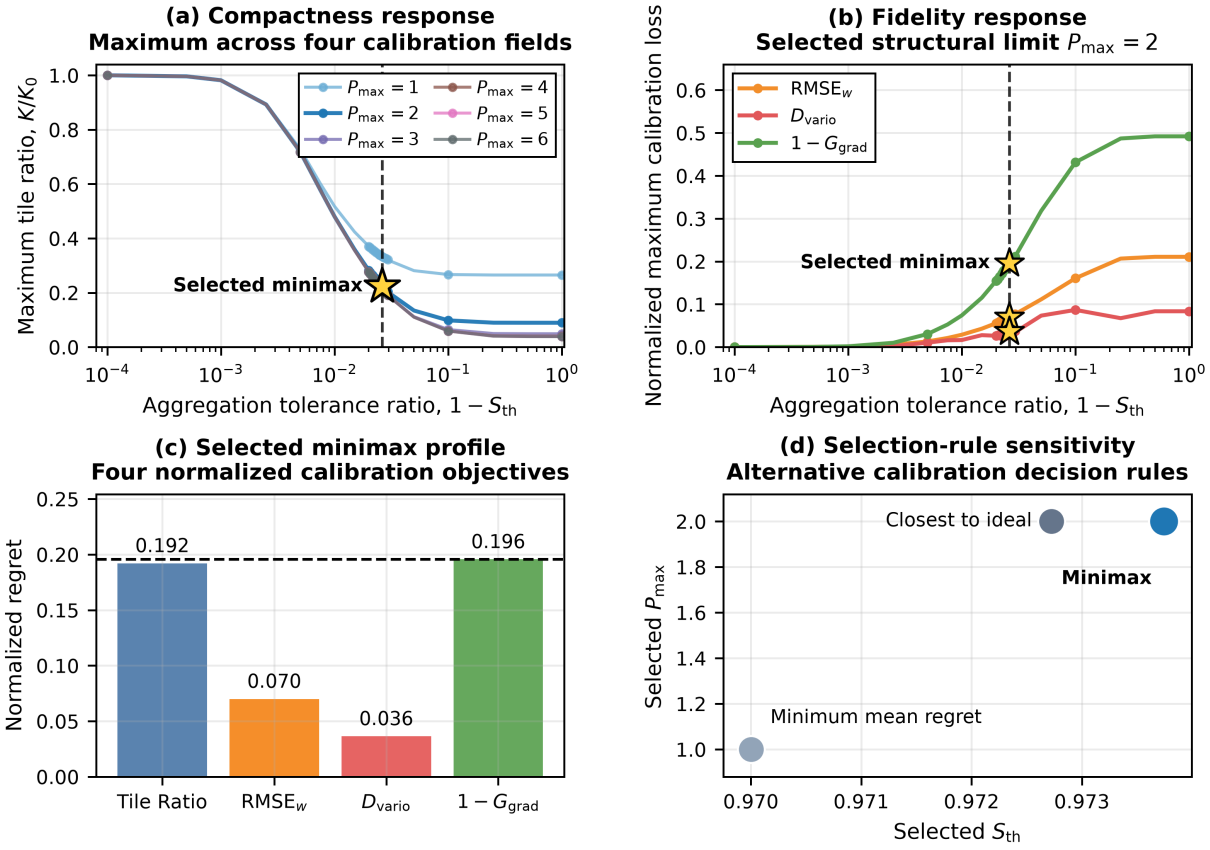


Figure 2. Parameter response and fixed-profile selection. (**a**) Structural tile ratio as a function of the aggregation-tolerance ratio $1 - S_{th}$ across maximum dyadic levels. The singleton case $P_{max} = 0$ is omitted, and $P_{max} = 7$ and 8 coincide with $P_{max} = 6$ throughout the displayed strict complete-block search, so only levels 1–6 are shown. (**b**) Normalized fidelity responses at the selected P_{max}. (**c**) Four normalized objective regrets of the selected calibration point. (**d**) Sensitivity of the selected operating point to alternative calibration-only decision rules. The star denotes $P_{max} = 2$ and $S_{th} = 0.9737439$, and the vertical dashed lines in (**a**,**b**) mark its horizontal-axis position; the held-out fields are not used in any panel. The horizontal dashed line in (**c**) marks the maximum normalized regret of the selected point.

3.2. Spatial Consequences of Candidate Selection

Figure 3 translates the scalar objectives into spatial outcomes. On M01, the selected adaptive grid used 27,513 tiles ($K/K_0 = 0.2114$), with RMSE$_w = 0.0835$, $D_{vario} = 0.0066$, and physical-gradient magnitude loss 0.1291. The nearest discrete uniform representation used more tiles ($K/K_0 = 0.2573$) yet produced much larger losses: RMSE 0.5949, normalized semivariogram discrepancy 0.0293, and physical-gradient magnitude loss 0.4763. Its 33,486 tiles comprised 379 full-valid 16×16 blocks and 33,107 singletons. The comparison therefore describes this strict-mask nearest-count rule, whose coarse interior blocks and singleton residuals can produce a poor fidelity balance; it does not establish an advantage over every fixed-block representation.

The compactness extreme ($P_{max} = 6$, $S_{th} = 0.25$) reduced the maximum calibration tile ratio to 0.0384 but increased the maximum calibration RMSE to 1.1932 and physical-gradient magnitude loss to 0.7276. Conversely, the candidate selected by minimum mean

regret ($P_{\max} = 1$, $S_{\text{th}} = 0.9700012$) used a maximum calibration tile ratio of 0.3188 and favored lower losses. The minimax profile lies between these states. Its fine tiles follow heterogeneous regions rather than a fixed shoreline-only template, while larger blocks occupy smoother interiors.

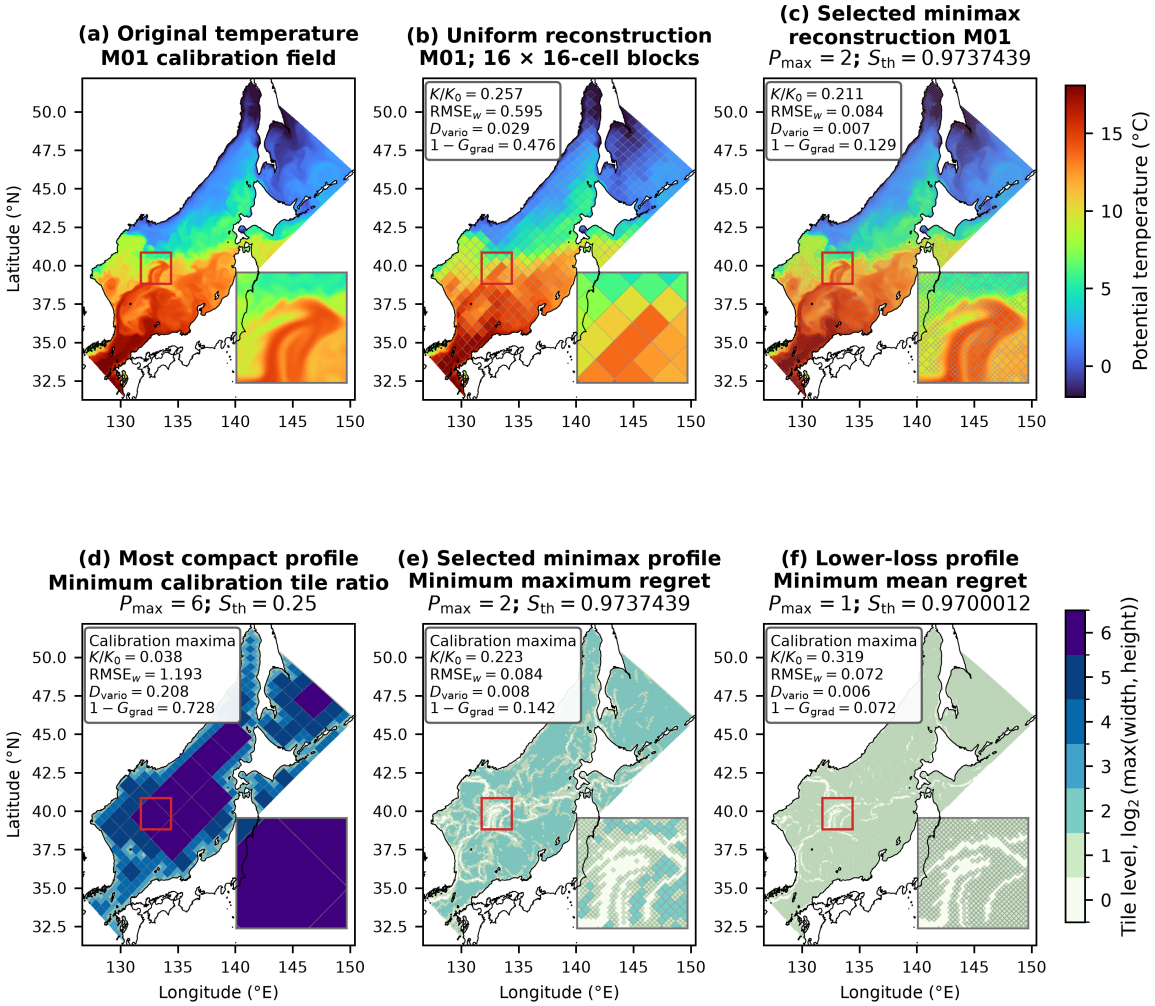


Figure 3. Spatial consequences of candidate selection. (**a**) Original M01 field. (**b**) Nearest discrete uniform representation for M01. (**c**) Selected minimax reconstruction for M01. (**d**) Compactness-prioritized calibration state. (**e**) Tile-level map of the selected calibration profile. (**f**) Candidate selected by minimum mean regret, which favors lower losses. Insets use a common narrower region to expose tile boundaries and local reconstruction behavior. Panels (**b**,**c**) report M01 values; panels (**d**–**f**) identify states using maxima over M01, M04, M07, and M10.

3.3. Spatial Comparison with Uniform and Quadtree Algorithms

The M05 comparison window in Figure 4 was selected from the held-out source field using physical-gradient energy only; no reconstruction metric determined its location. This deliberately structured window is a spatial diagnostic rather than a representative random sample. The 80 × 120-cell window has an ocean-valid fraction of 0.8541. The proposed representation had local RMSE 0.0687, normalized semivariogram discrepancy 0.0054, and physical-gradient magnitude loss 0.0710, compared with 0.6010, 0.1618, and 0.5620 for the uniform representation. Visually, the uniform result spreads error across structured features, whereas adaptive methods retain smaller tiles around locally varying regions.

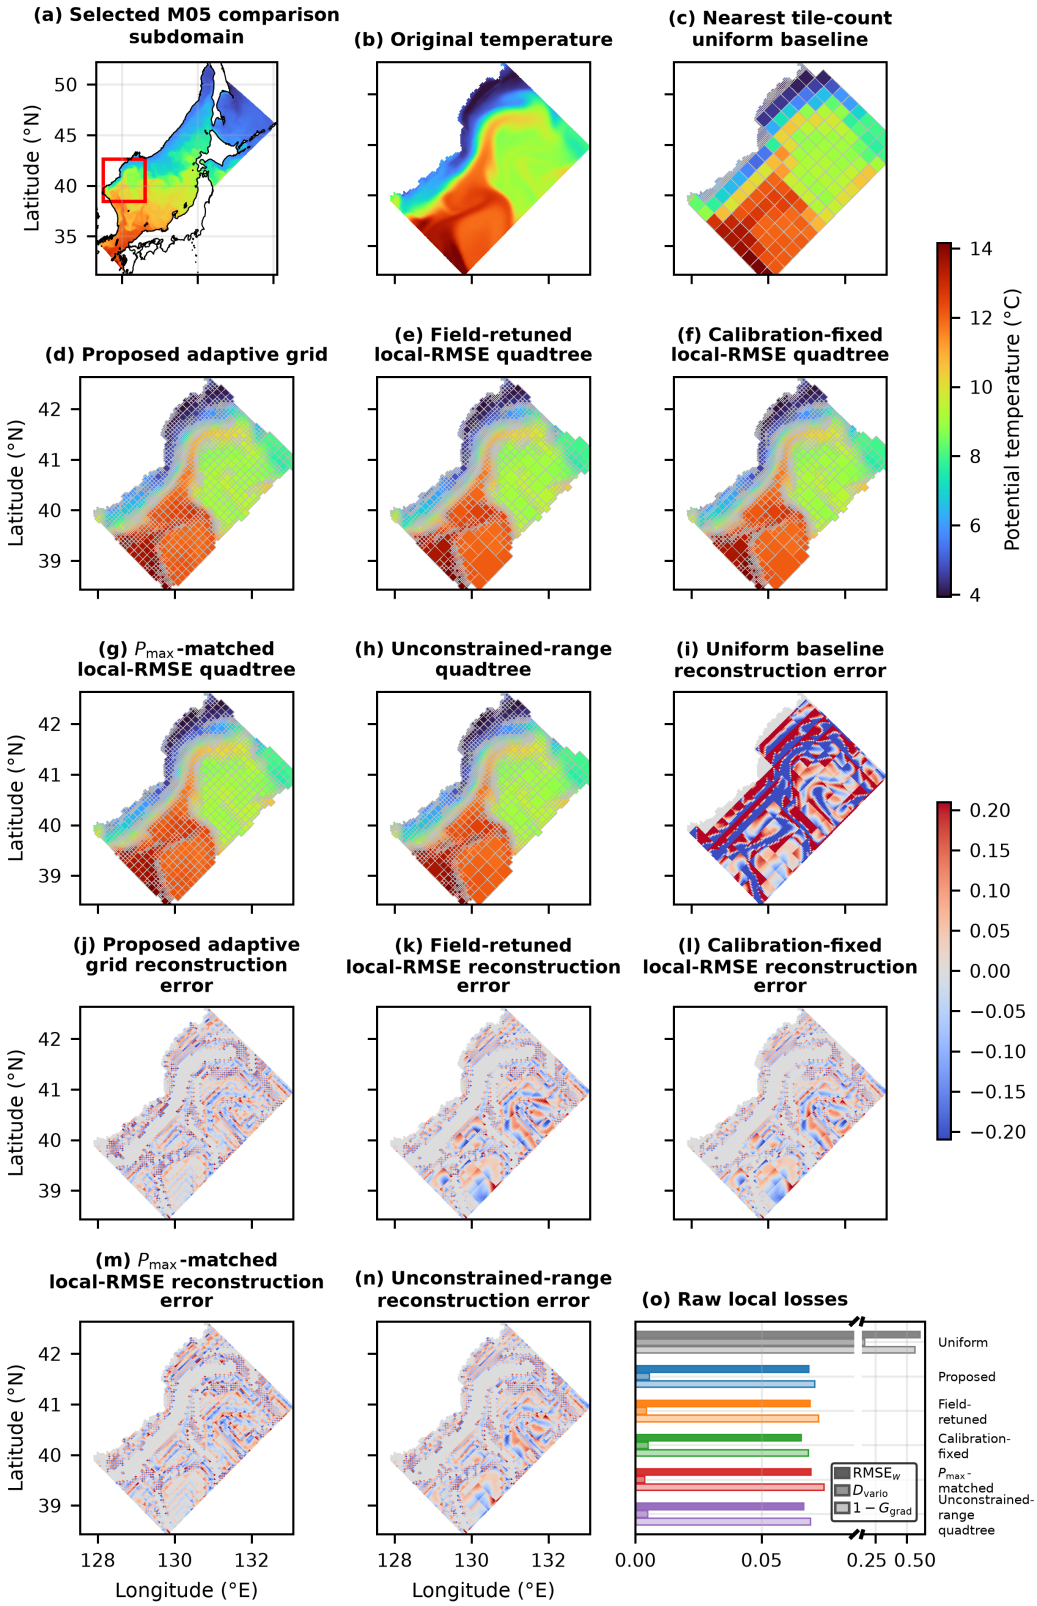


Figure 4. Spatial comparison on M05. (**a**) Comparison subdomain; (**b**) original field; (**c**) uniform representation; (**d**) proposed adaptive grid; and (**e**–**h**) quadtree representations. (**i**–**n**) Reconstruction errors on a common color scale. (**o**) Raw local RMSE, semivariogram discrepancy, and gradient-magnitude loss. M05 is held out, and the displayed subdomain is diagnostic only.

No adaptive construction ranks first across all three local metrics. Field-retuned and P_{max}-matched local-RMSE variants address rate matching after field-specific adjustment, whereas calibration-fixed methods test transfer. We therefore interpret the proposed method as a balanced fixed-profile representation, not the local winner; the common error-map scale enables direct spatial comparison.

All methods in Figure 4 use the same source cells, validity mask, geographic window, and displayed color limits. Consequently, the differences among the reconstruction and residual panels arise from how each partition allocates support, not from a change in crop or evaluation domain. The original field and the reconstructed fields are shown separately from the residual maps so that a visually smooth reconstruction is not mistaken for a small error. The uniform baseline illustrates this distinction most clearly: its broad blocks obscure organized temperature variation, and the corresponding residuals remain spatially structured rather than resembling uncorrelated fine-scale noise.

The four adaptive comparators form a much narrower local-loss range than the uniform baseline. Across those methods and the proposed adaptive grid, local RMSE spans 0.0656–0.0695, D_{vario} spans 0.0037–0.0054, and physical-gradient magnitude loss spans 0.0685–0.0748. The calibration-fixed local-RMSE quadtree has the smallest RMSE and gradient loss in this window, whereas the P_{max}-matched variant has the smallest variogram discrepancy. These reversals show why a single locally selected window cannot establish an overall method ranking and why the three fidelity measures are retained rather than collapsed into one post hoc score.

The partition and error panels provide information that the raw values in panel (o) cannot convey alone. Similar loss values may accompany different tile boundaries and therefore different neighborhoods for spatial queries, selective transfer, or subsequent aggregation. Conversely, visibly different partitions need not imply a practically large change in every field-scale metric. Figure 4 is thus used as a spatial diagnostic of allocation behavior; the fixed-profile held-out summaries provide the quantitative cross-field comparison.

3.4. Calibration Operating Sets Across Spatial Methods

Figure 5 distinguishes two levels of comparison. The upper row shows shared-parameter candidate sets generated during calibration, whereas the lower row shows the six calibration policies that are subsequently evaluated field by field. In every column, the horizontal coordinate is the maximum tile ratio across M01, M04, M07, and M10, and the vertical coordinate is the corresponding maximum fidelity loss. Thus, movement toward the lower left is preferable within a panel because it simultaneously indicates fewer tiles and a smaller loss. This visual direction does not, however, establish universal dominance: the three loss metrics emphasize different aspects of fidelity, and some policy points are obtained under different rate-matching rules.

The lower row contains the proposed adaptive grid, nearest tile-count uniform baseline, Field-retuned local-RMSE quadtree, calibration-fixed local-RMSE quadtree, proposed-depth rate-matched local-RMSE quadtree, and unconstrained-range quadtree. The final label denotes the Unconstrained-range comparator in Table 2. The proposed profile is highlighted by a black-outlined blue star with a translucent fill, drawn above the comparator markers. In the upper row, the two local-RMSE candidate sets each use one normalized threshold shared across all four calibration fields: the structural-ceiling sweep fixes $P_{max} = 8$, the prespecified structural ceiling, whereas the proposed-depth sweep fixes $P_{max} = 2$, the depth selected for the proposed profile. These operating sets must not be confused with the Field-retuned and proposed-depth rate-matched points in the lower row, whose tolerances are adjusted separately for each field to approach the proposed tile count. The proposed-depth rate-matched point is the same P_{max}-matched comparator

defined in Table 2; its shorter figure label emphasizes the shared depth without adding another symbol to the legend. Normalized regret is defined within each candidate family for sensitivity analysis and is not compared across method families.

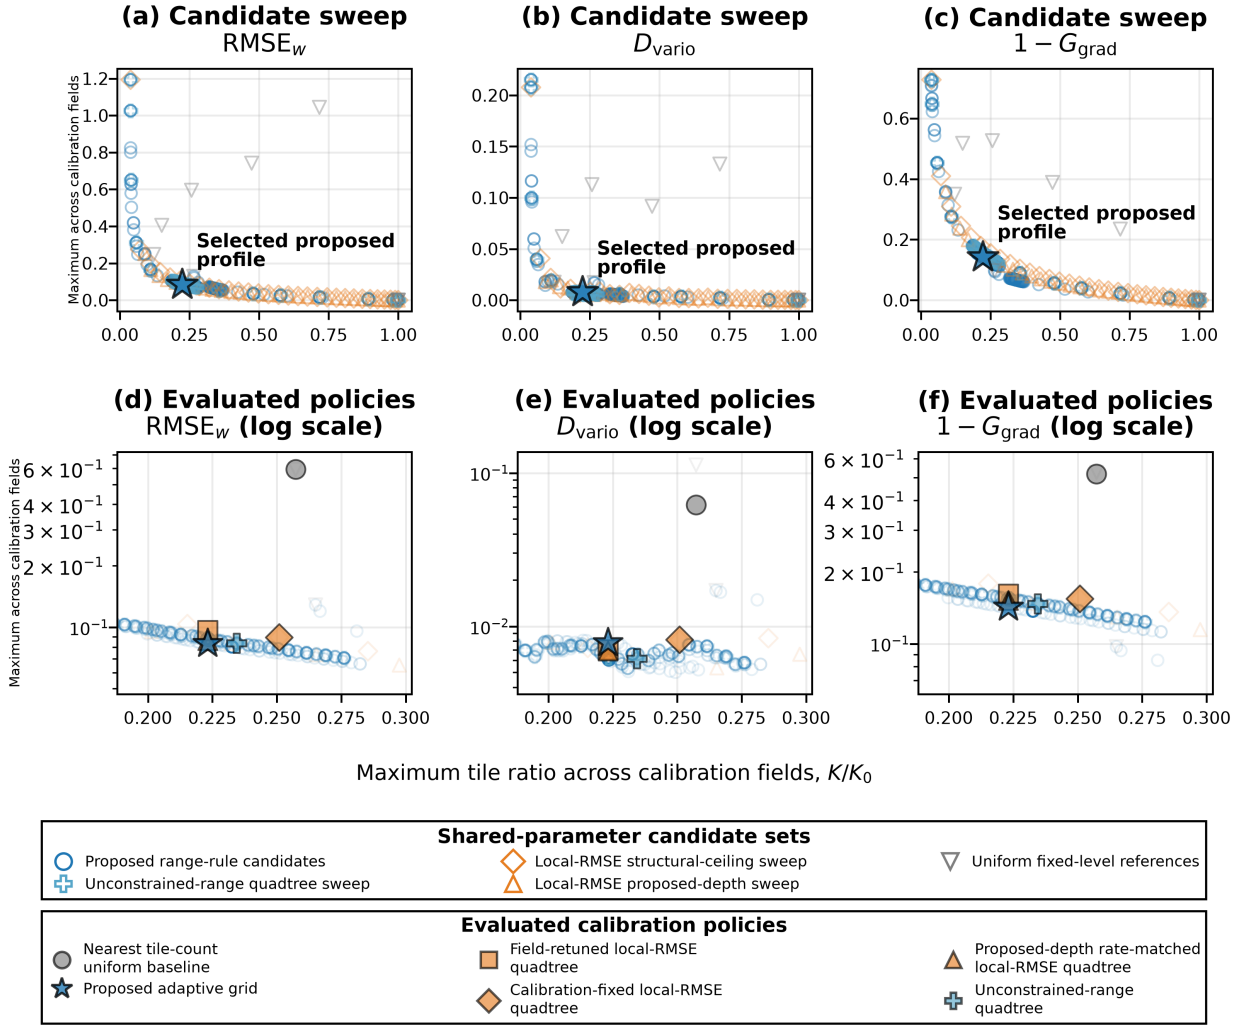


Figure 5. Calibration candidate sets and evaluated policies. Columns compare the maximum tile ratio across the four calibration fields with maximum (**a**,**d**) RMSE$_w$, (**b**,**e**) D_{vario}, and (**c**,**f**) $1 - G_{\text{grad}}$. Lower-left locations are preferable within each panel because both structural burden and fidelity loss are smaller. Open markers in the upper row show shared-parameter candidate sweeps. In the lower row, translucent filled markers identify the five comparators, whereas the black-outlined blue star highlights the proposed four-objective minimax profile. Blue denotes range-rule constructions, orange denotes local-RMSE constructions, and gray denotes uniform references. The Field-retuned and proposed-depth rate-matched local-RMSE policies use field-specific tolerances and therefore appear as evaluated points rather than shared-parameter calibration curves. The unconstrained-range policy is the Unconstrained-range comparator in Table 2. The lower-row loss axes use logarithmic scaling to keep all evaluated policies legible.

The two rows therefore answer different questions. A point in the upper row describes one candidate generated under a parameter-sharing rule on the calibration subset; it does not yet represent a held-out evaluation policy. A symbol in the lower row identifies the fully specified policy that is carried into the subsequent comparisons, including whether its tolerance is fixed from calibration or adjusted for each field. Keeping these layers separate prevents a field-retuned rate match from being interpreted as evidence of parameter transfer.

The apparent traces in the upper row should also be read as sequences of discrete realized partitions, not as smooth response functions. Adjacent threshold values can produce the same tile layout, whereas a small threshold change at a merge boundary can move several objectives simultaneously. Open and translucent markers preserve this candidate structure without implying interpolation between unattained states. The shared horizontal coordinate makes structural burden comparable, while the three vertical coordinates expose the metric-specific consequences of a candidate.

In the lower row, policies with broadly similar calibration tile ratios remain separated in at least one fidelity panel, and the nearest uniform baseline lies well above the adaptive cluster for some losses. The unconstrained-range policy is close to the proposed policy in aggregate calibration space, but it is retained as a separately calibrated comparator because its maximum level and threshold differ and it omits the physical safeguard. Accordingly, Figure 5 motivates the held-out analysis rather than deciding the comparison by itself: a favorable lower-left position is necessary for a two-objective view, but consistent behavior across fields and metrics must be checked independently.

3.5. East Sea Within-Sequence Held-Out Evaluation

The fixed proposed profile retained a median 0.2049 of the valid-cell count as tiles across the eight held-out fields (95% field-bootstrap interval 0.1977–0.2289). Its median RMSE was 0.0708 (0.0653–0.0760), median normalized semivariogram discrepancy was 0.0051 (0.0029–0.0059), and median physical-gradient magnitude loss was 0.1304 (0.1210–0.1389). Figure 6 and Table 3 report all six methods using the same fields and summaries.

The nearest-count uniform representation had a similar median tile ratio but substantially greater median losses. Relative to the calibration-fixed local-RMSE quadtree, paired proposed-minus-comparator intervals were below zero for RMSE (-0.0058 to -0.0035) and physical-gradient magnitude loss (-0.0114 to -0.0050). The normalized-semivariogram interval was -0.0025 to 0.0001 and therefore included zero. These intervals describe the eight sampled fields rather than independent-population superiority. The proposed method also had lower held-out losses than the Field-retuned local-RMSE quadtree at a matched median tile ratio. A local RMSE split statistic therefore need not minimize the resulting field-scale losses, even though the calibration-fixed families share the same multiobjective selection convention.

Fixed-block references qualify the nearest-count comparison. The strict-mask 2×2 reference had median $K/K_0 = 0.2650$, $\mathrm{RMSE}_w = 0.1095$, $D_{\mathrm{vario}} = 0.0087$, and gradient loss 0.0949. The proposed profile therefore used fewer tiles and reduced RMSE and semivariogram discrepancy, but the 2×2 reference retained gradient magnitude more accurately. The strict-mask 4×4 reference used fewer tiles ($K/K_0 = 0.1215$) with larger losses (0.2153, 0.0105, and 0.3327, respectively). Mask-partial fixed-block references, which average only valid cells within each occupied footprint, show the same qualitative trade-off. Field-level and summary values for both mask policies accompany the Supplementary Data. These references prevent the particular nearest-count rule from being interpreted as a universal uniform-grid benchmark.

The unconstrained-range quadtree is an important qualification. Its median losses were marginally lower than those of the proposed profile, and the paired intervals for all three losses included zero. However, it is not a true ablation: its maximum level and calibration-selected threshold both differ from those of the proposed profile. Its held-out median cell partition-change fraction relative to the proposed representation was 0.3078, despite similar aggregate metrics. The result supports a narrower novelty claim: the

contribution is a deterministic, constraint-based representation and calibration protocol with exact membership, not a demonstrated fidelity gain from a new split heuristic.

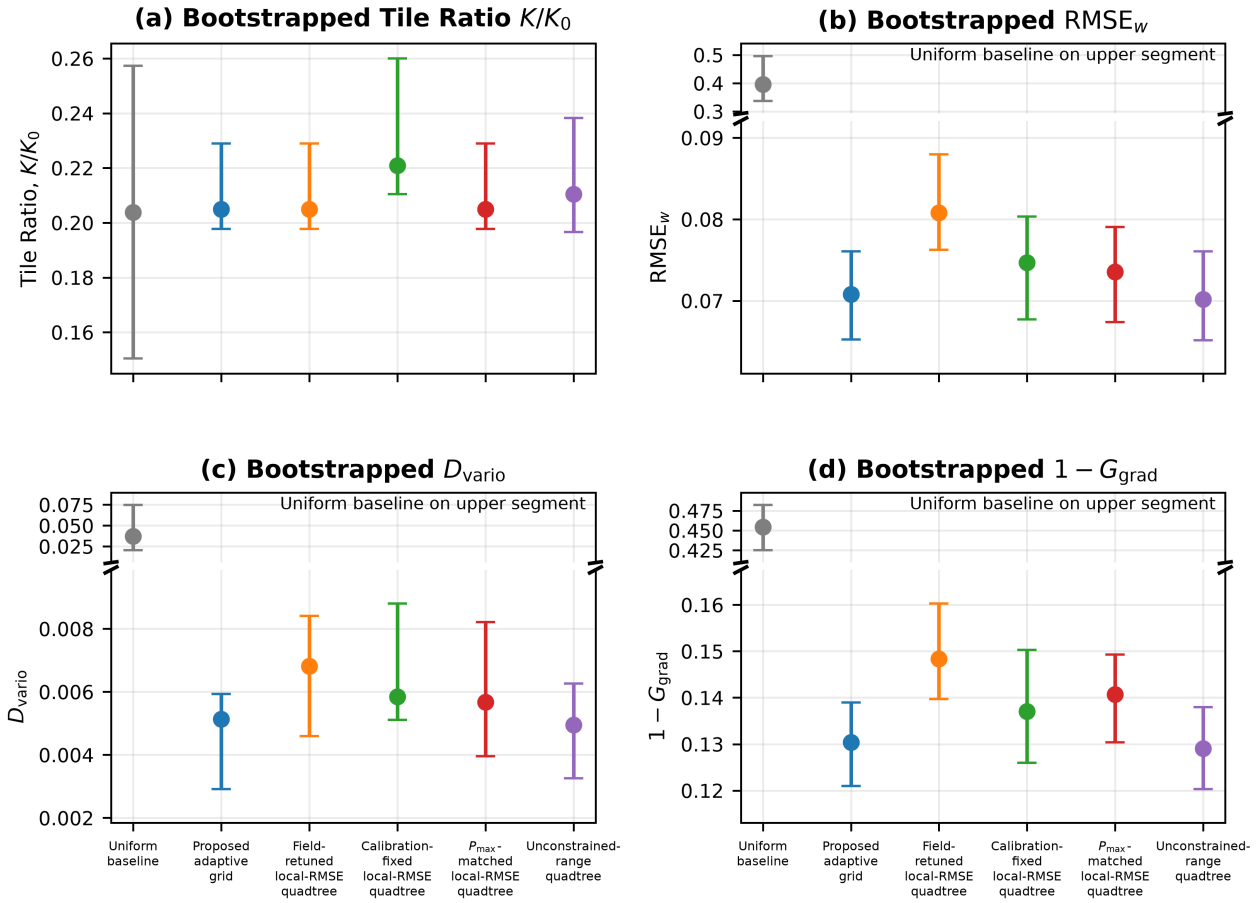


Figure 6. East Sea within-sequence held-out performance of the six spatial methods. Points are medians across the eight noncalibration nominal monthly first-day fields; bars are 2.5th–97.5th percentiles from 5000 field-level bootstrap resamples. (**a**) Structural tile ratio K/K_0. (**b**) Area-weighted RMSE. (**c**) Normalized semivariogram discrepancy. (**d**) Physical-gradient magnitude loss. A partition is regenerated from each field. Broken axes enlarge the adaptive-method region where the uniform result is much larger.

Table 3. Held-out medians with 95% field-bootstrap intervals for spatial methods. Lower is better for all four columns. Intervals summarize eight held-out fields and do not imply independent population sampling.

Method	K/K_0	RMSE_w	D_{vario}	$1 - G_{\mathrm{grad}}$
Proposed adaptive grid	0.2049 [0.1977, 0.2289]	0.0708 [0.0653, 0.0760]	0.0051 [0.0029, 0.0059]	0.1304 [0.1210, 0.1389]
Nearest tile-count uniform baseline	0.2038 [0.1504, 0.2573]	0.3948 [0.3363, 0.4954]	0.0373 [0.0205, 0.0745]	0.4543 [0.4252, 0.4819]
Field-retuned local-RMSE quadtree	0.2049 [0.1977, 0.2290]	0.0808 [0.0763, 0.0880]	0.0068 [0.0046, 0.0084]	0.1484 [0.1397, 0.1603]
Calibration-fixed local-RMSE quadtree	0.2208 [0.2104, 0.2600]	0.0747 [0.0677, 0.0803]	0.0058 [0.0051, 0.0088]	0.1370 [0.1260, 0.1503]
$P_{\max}$-matched local-RMSE quadtree	0.2049 [0.1977, 0.2289]	0.0735 [0.0674, 0.0791]	0.0057 [0.0040, 0.0082]	0.1407 [0.1304, 0.1493]
Unconstrained-range quadtree	0.2104 [0.1966, 0.2383]	0.0702 [0.0651, 0.0761]	0.0049 [0.0033, 0.0063]	0.1291 [0.1203, 0.1380]

Figure 6 and Table 3 provide complementary views of the same held-out summaries. The broken axes enlarge the adaptive-method region without changing any reported value, while the table retains the medians and field-bootstrap intervals. Similar tile ratios should therefore be interpreted jointly with all three fidelity losses and not as evidence of equivalent partitions.

3.6. Exact Realized-Tile Scale-Cap Sensitivity

The four East Sea temperature calibration ranges were 701.28, 524.20, 737.76, and 891.52 km, so Equation (6) gave a fixed cap of 366.9376 km. This cap blocked no merges across the evaluated primary candidate set. At the selected profile, the maximum realized tile diameter was 21.157 km on the calibration fields and 21.293 km across all twelve fields. No fixed-profile merge was blocked at 30 km or at any larger tested cap, and the partitions at the correlation-derived cap and 400 km were identical. The cap therefore serves as a formal safeguard in this experiment; the primary results do not establish an independent performance contribution from it.

Figure 7 separates two analysis protocols. Panels (a,b,f) apply the fixed calibration-selected profile at each test cap. At 20 km, the cap blocked at least one merge in all calibration (4/4) and held-out (8/8) fields and changed the containing tile relative to the uncapped fixed profile for median fractions of 0.0559 and 0.0616 of their valid cells, respectively. At 10 km, the corresponding topology-change fractions were 0.6425 and 0.6593, again with all fields binding. Caps of 30 km and greater produced zero topology change and no binding fields. The 20 km fixed-profile cap raised the maximum calibration tile ratio from 0.2231 to 0.2341 while reducing the maximum calibration RMSE from 0.0835 to 0.0817 and gradient-magnitude loss from 0.1424 to 0.1349; the 10 km values were 0.4030, 0.0605, and 0.0592.

Panels (c–e) instead recalibrate the profile separately at each cap using only the four calibration fields. The selected maximum level was one at 10–30 km and two at 50 km and above. Under this cap-specific calibration protocol, only the 10 km profile bound any fields (4/4 calibration and 8/8 held out); the 20 km and larger cap-specific profiles remained below their caps. These results demonstrate the constraint's effect when active without attributing the primary fixed-profile performance to an inactive cap.

The fixed-profile and cap-specific views answer different questions and should not be combined into a single response curve. Holding P_{max} and S_{th} fixed isolates the cap's mechanical effect on an otherwise unchanged profile, whereas recalibration allows the other controls to adapt to each cap. The binding count verifies whether the constraint is encountered, and the partition-change fraction measures the resulting topological effect. Thus, the primary cap remains part of the admissible method definition, but the East Sea results cannot be attributed to it because it is inactive there. Conversely, the 10 and 20 km perturbations verify implementation and control when the constraint binds; they are diagnostic tests, not recommended physical settings.

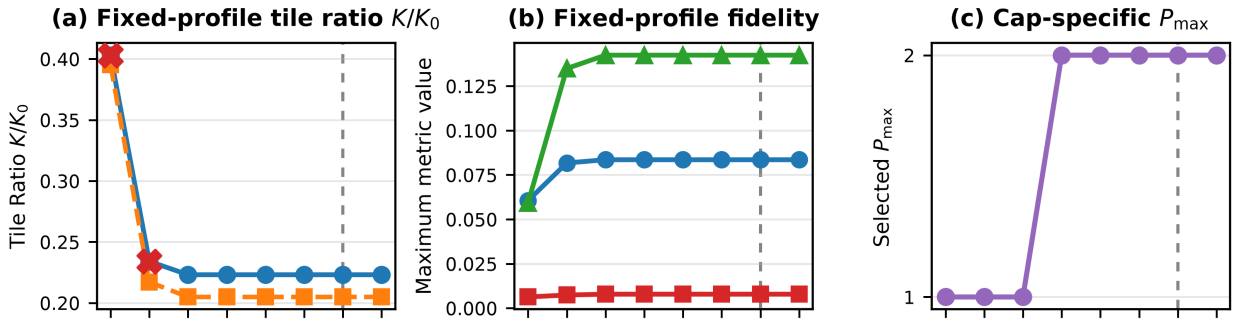


Figure 7. *Cont.*

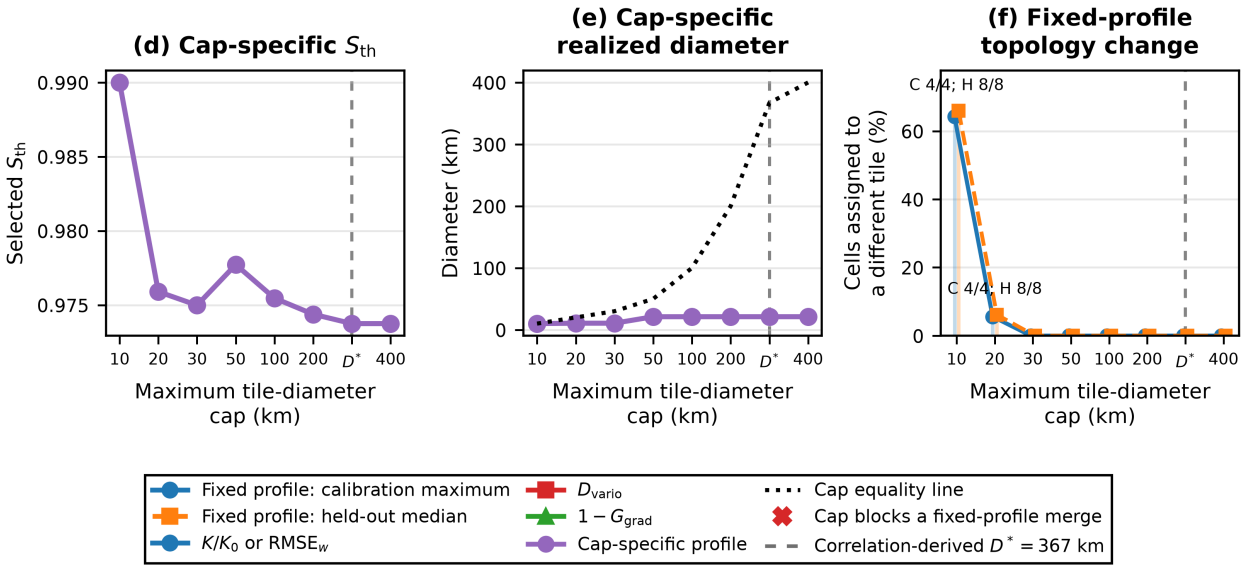


Figure 7. Exact realized-tile scale-cap sensitivity under two explicitly separated analysis protocols. Panels (**a**,**b**,**f**) apply the fixed calibration-selected temperature profile at every cap: (**a**) calibration-maximum and held-out-median tile ratios; (**b**) calibration-maximum fidelity losses; and (**f**) the median fraction of valid cells whose containing tile differs from the uncapped fixed-profile partition, shown separately for calibration and held-out fields. This fixed profile binds in 4/4 calibration and 8/8 held-out fields at 10 and 20 km, and in 0 fields at caps of 30 km or greater. Panels (**c**–**e**) use cap-specific calibration: (**c**) selected P_{max}, (**d**) selected S_{th}, and (**e**) maximum realized tile diameter. Only the 10 km cap-specific profile binds (4/4 calibration and 8/8 held-out fields). The vertical dashed line marks the 366.9376 km correlation-derived cap. Tile diameter is the maximum WGS84 geodesic distance among the six pairs of outer block corners.

3.7. Cross-Variable and External-Observation Validation Results

Figure 8 and Table 4 separate direct temperature-profile transfer without recalibration from target-specific calibration. On East Sea salinity, applying the complete temperature profile without recalibration produced a held-out median $K/K_0 = 0.2298$, RMSE$_w/R = 0.0045$, $D_{vario} = 0.0021$, and physical-gradient magnitude loss 0.1095. The salinity-specific minimax profile selected $P_{max} = 1$ and $S_{th} = 0.380145$, with a larger median tile ratio (0.2651) and normalized RMSE (0.0093) but lower physical-gradient magnitude loss (0.0926).

The primary VIIRS analysis selected $P_{max} = 1$ and $S_{th} = 0.8875$ from January, April, July, and October composites of the dimensionless field $\log_{10}[C/(1\ \mathrm{mg\,m^{-3}})]$. Across the eight held-out composites, its median K/K_0 was 0.4055, with RMSE$_w/R = 0.0148$, $D_{vario} = 0.0068$, and physical-gradient magnitude loss 0.1744. Direct application of the temperature profile retained substantially more tiles (median $K/K_0 = 0.7422$) and lower losses, demonstrating the compactness–fidelity shift that occurs when a profile selected for another variable, grid, region, and source is transferred without recalibration. The raw-concentration and all-12-common-support sensitivities likewise show that transform, support, and calibration definitions affect the operating point.

The scale screen was nonbinding in both validation case studies. Two of four salinity calibration semivariogram fits and all four VIIRS calibration fits reached the permitted range upper bound. Accordingly, salinity and chlorophyll-a broaden the representation test but supply no additional empirical evidence that the correlation-derived physical cap is active; in particular, the external VIIRS range is not treated as a resolved correlation length.

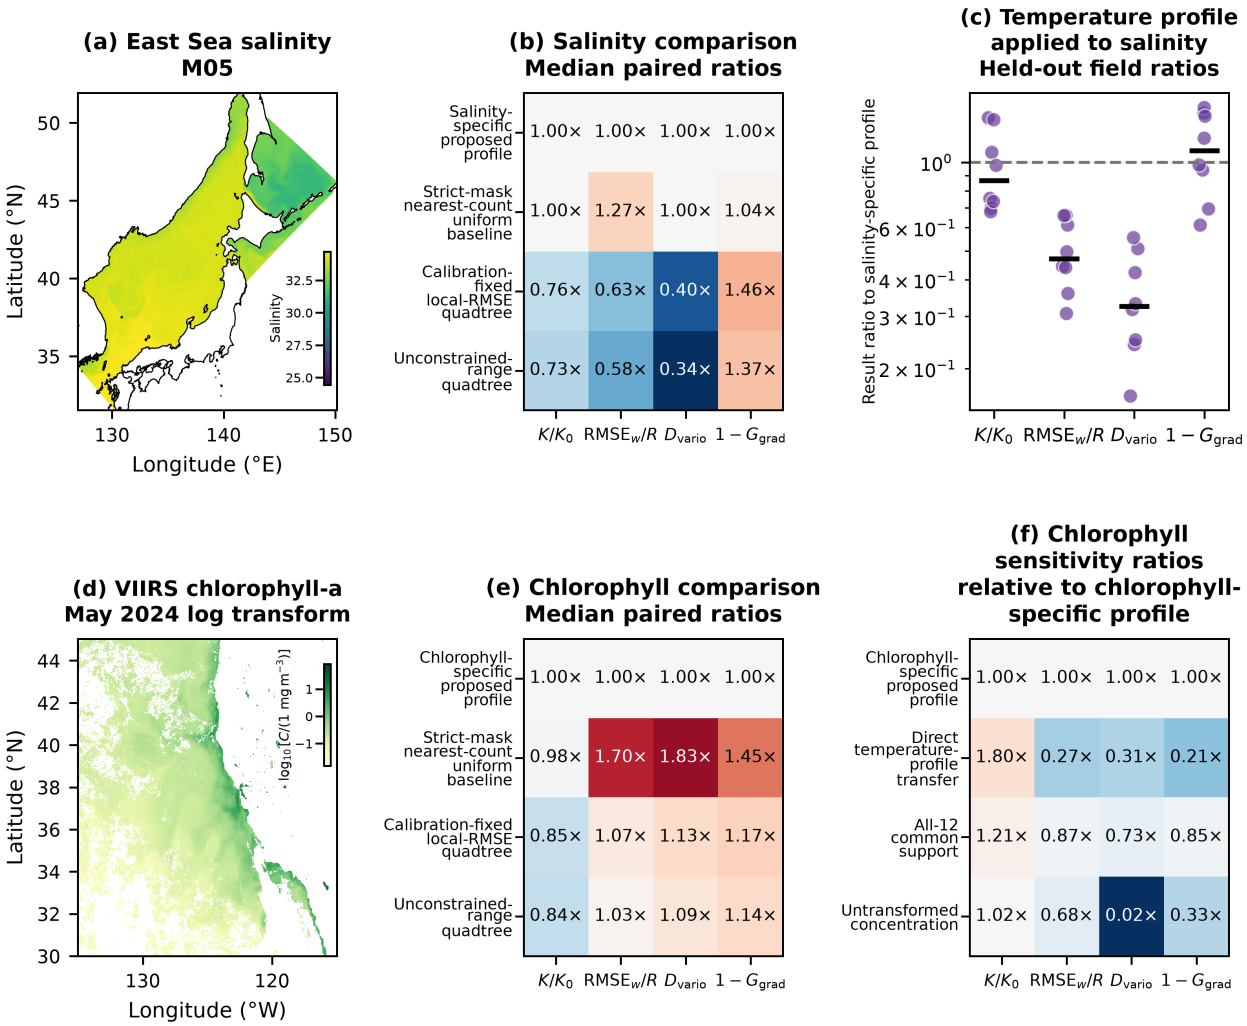


Figure 8. Cross-variable and external-observation validation. (**a**) East Sea M05 near-surface salinity. (**b**) Median field-paired held-out salinity ratios relative to the salinity-specific proposed profile for structural tile ratio, normalized RMSE, normalized semivariogram discrepancy, and physical-gradient magnitude loss. (**c**) Field-level ratios obtained by direct temperature-profile transfer to salinity without recalibration. (**d**) California Current NOAA VIIRS chlorophyll-a for May 2024, displayed and analyzed primarily as the dimensionless field $\log_{10}[C/(1\ \mathrm{mg\,m^{-3}})]$. (**e**) Median field-paired held-out VIIRS ratios relative to the chlorophyll-specific proposed profile. (**f**) Chlorophyll support, transform, and direct-transfer sensitivities. Relative summaries are formed within each field before aggregation. Raw-concentration versus log-transformed ratios in (**f**) describe transform sensitivity rather than errors on a common physical scale. The scale screen was nonbinding in these case studies. In (**b**,**e**,**f**), blue and red indicate ratios below and above 1, respectively, with white centered at 1; color intensity follows a panel-specific $\log_2$ ratio scale. In (**c**), purple dots show individual held-out field ratios, black horizontal bars show their medians, and the dashed line marks a ratio of 1.

Panels (b,e) compare methods within the same target dataset, using calibration-fixed adaptive profiles and the field-specific nearest-count uniform rule, whereas panels (c,f) isolate profile transfer and analysis choices. This distinction matters because retaining more tiles can lower a loss without demonstrating successful parameter transfer. The two case studies therefore support portability of the deterministic construction and evaluation protocol, but not a universal operating profile. Their differing compactness–fidelity balances also motivate reporting every metric alongside the selected parameter pair.

Table 4. Cross-variable and external-observation held-out medians. RMSE_w/R is dimensionless and supports within-table comparison more appropriately than raw variable-specific RMSE. Each row summarizes eight noncalibration fields. A profile is reported as $(P_{\max}, S_{\mathrm{th}}, D_{\mathrm{cap}})$; target-specific profiles use only their four calibration fields.

Dataset	Profile Definition	K/K_0	RMSE_w/R	D_{vario}	$1 - G_{\mathrm{grad}}$
East Sea salinity	Direct temperature-profile transfer without recalibration: $(2, 0.9737439, 366.9376\ \mathrm{km})$	0.2298	0.0045	0.0021	0.1095
East Sea salinity	Salinity-specific proposed profile: $(1, 0.380145, 1084.3696\ \mathrm{km})$	0.2651	0.0093	0.0063	0.0926
California Current VIIRS transformed chlorophyll-a	Direct temperature-profile transfer without recalibration: $(2, 0.9737439, 366.9376\ \mathrm{km})$	0.7422	0.0040	0.0017	0.0380
California Current VIIRS transformed chlorophyll-a	Chlorophyll-specific proposed profile: $(1, 0.8875, 1744.5918\ \mathrm{km})$	0.4055	0.0148	0.0068	0.1744

3.8. Monthly First-Day Topology Variation

Across all twelve fields, the mean tile ratio was 0.2126 and ranged from 0.1919 to 0.2600 (Figure 9c). A broadly stable representation burden therefore did not imply a fixed partition. The mean fraction of cells whose exact containing tile rectangle changed between adjacent fields was 0.3649 (range 0.2989–0.4468), while exact-rectangle Jaccard similarity averaged 0.3310 (range 0.2715–0.3766). The mean cellwise persistence of the modal tile level was 0.7100.

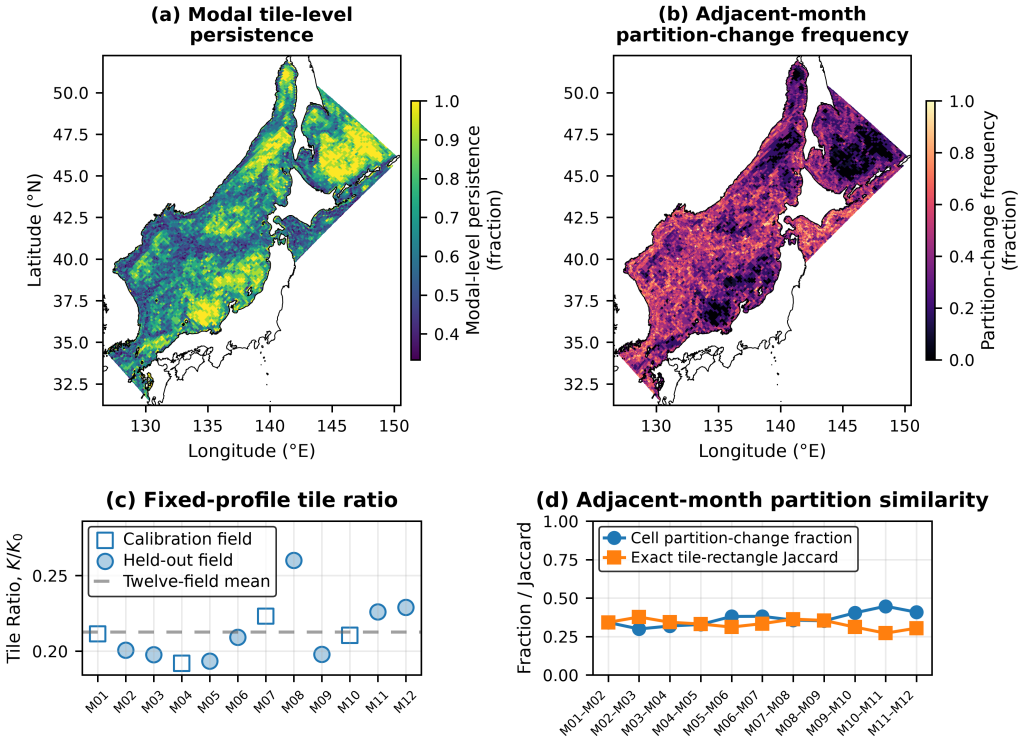


Figure 9. Monthly first-day adaptive-partition topology diagnostic. (**a**) Persistence of each cell's modal tile level. (**b**) Frequency with which the exact containing tile rectangle changes across 11 adjacent monthly first-day pairs. (**c**) Monthly tile ratio; the translucent gray dashed line is the twelve-field mean. (**d**) Adjacent-pair cell partition-change fraction and exact tile-rectangle Jaccard similarity. Similar structural burden can coexist with relocation of exact tile boundaries. The diagnostic does not resolve daily or within-month variation.

These results separate two notions of monthly first-day consistency. Approximate scale assignment was persistent over much of the domain, but exact refinement locations

moved with the field. Because only one nominal field per month is analyzed, these results do not characterize daily, synoptic, or within-month topology changes and do not establish flicker-free performance at high-frequency operational time steps. Applications requiring temporally coherent topology would need denser sampling plus an explicit temporal regularizer, persistent hierarchy, or update-cost constraint.

3.9. Fixed-Budget Model-Mask-Boundary Allocation Results

Figure 10 tests whether the fixed partition rule can redirect a constant tile budget toward a prescribed region. The numerical boundary is the in-domain valid–invalid interface of the M01 source-model mask; Natural Earth linework is included only for geographic context. For each positive strength ρ, the multiplier c_ρ is chosen on M01 to recover the global-reference count $K_{\mathrm{ref}} = 27{,}513$ exactly. At $\rho = 0.60$, this requires $c_\rho = 1.2691$.

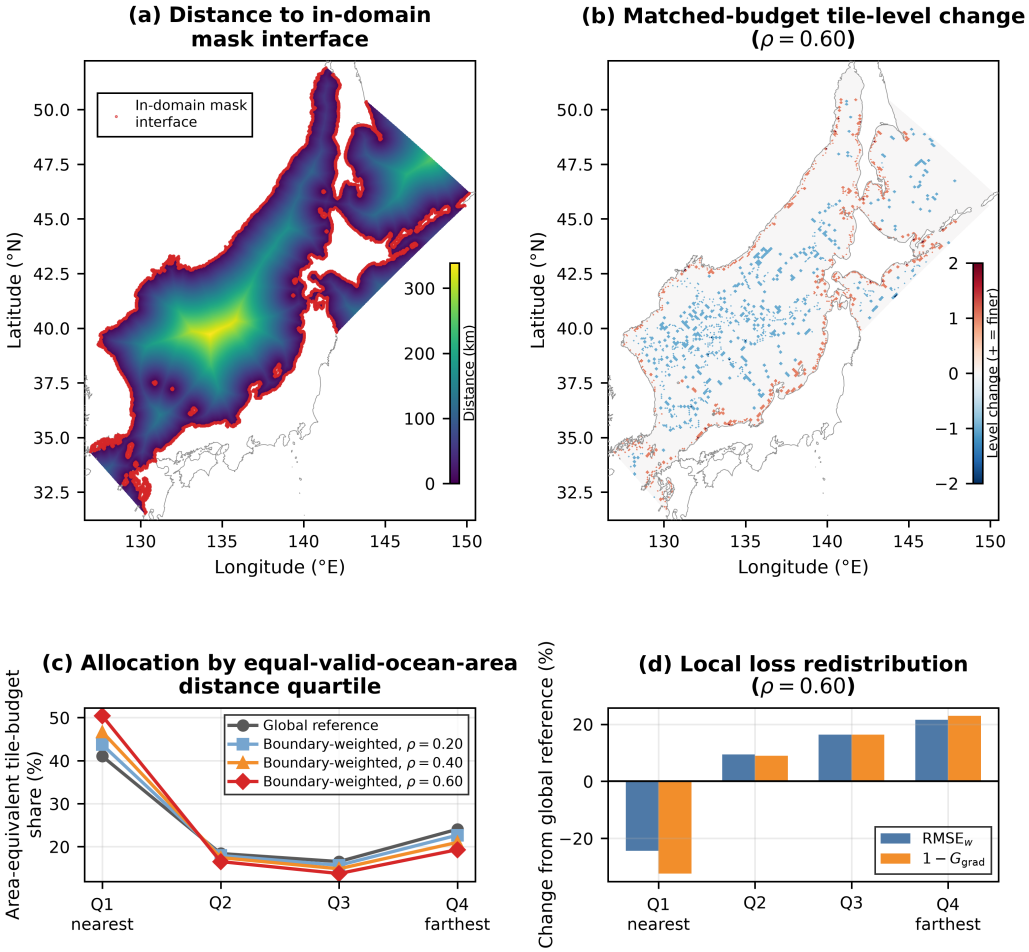


Figure 10. Fixed-budget model-mask-boundary allocation on M01. (**a**) AEQD distance from each valid cell to the nearest in-domain valid–invalid mask-interface cell; Natural Earth linework is visual context only. (**b**) Change in tile level for the $\rho = 0.60$ matched-budget partition relative to the global reference, with positive values denoting finer allocation. (**c**) Area-equivalent share of the fixed tile budget assigned to four equal-valid-ocean-area distance quartiles for $\rho = 0$, 0.20, 0.40, and 0.60. (**d**) Percentage changes in quartile-local area-weighted RMSE and physical-gradient magnitude loss at $\rho = 0.60$ relative to the global reference. Negative values indicate lower local loss. All matched-budget cases use $K = 27{,}513$; the same-base-tolerance ($c_\rho = 1$) cost diagnostic is reported separately in Supplementary Table S7.

As ρ increases from 0 to 0.60, the area-equivalent share assigned to the nearest boundary-distance quartile rises from 0.4102 to 0.5042. At $\rho = 0.60$, local RMSE and gradient loss in that nearest quartile fall by 24.37% and 32.30%, respectively, whereas the corresponding losses in the farthest quartile rise by 21.59% and 22.96%. The redistribution also changes global metrics: RMSE increases from 0.0835 to 0.0891 and gradient loss from 0.1291 to 0.1339, while normalized semivariogram discrepancy decreases from 0.0066 to 0.0052. Thus the fixed-budget mechanism prioritizes the specified interface neighborhood by shifting representation away from distant cells; it does not improve every region or every global metric.

The separate $\rho = 0.60$, $c_\rho = 1$ diagnostic retains the global base tolerance rather than the tile budget. It uses 32,394 tiles and yields global RMSE 0.0739, normalized semivariogram discrepancy 0.0051, and gradient loss 0.1093. These lower losses are accompanied by a 17.74% increase in tile count and therefore are not evidence of a free improvement. Supplementary Table S7 records all fixed-budget strengths and this same-base-tolerance cost diagnostic.

4. Discussion

4.1. Interpretation as a Derived Spatial Layer

The proposed method is a deterministic spatial summarization layer whose output comprises both a reconstructed array and an explicit partition. Each tile has a location, dyadic level, representative value, and exact membership relation to the native valid cells. This structure is intended to support regional selection and visual explanation of where a field requires finer representation. Level-of-detail display and selective transfer are potential applications; the current flat tile-record format does not itself implement progressive multiresolution decoding. Task completion, query latency, and user performance were not evaluated.

The distinction from a model grid is substantive. ROMS, adaptive ocean solvers, the Finite-volumE Sea ice–Ocean Model (FESOM2), and the ocean component of the Model for Prediction Across Scales (MPAS-Ocean) use discretizations whose geometry participates in governing-equation operators, fluxes, conservation, and coupling [21–23,37]. Our tiles are generated after the calculation and are piecewise-constant. They must not be used as though they preserved staggered-grid velocities, divergence, vorticity, pressure gradients, or energetic consistency. A consumer requiring those diagnostics should access the source field and model-grid metadata.

Within that boundary, the representation addresses needs identified in marine feature analysis and digital-twin systems: a compact visual layer can expose heterogeneous regions while retaining an exact tile-to-cell membership map to authoritative data [4–6,8]. Its explicit spatial map is also directly inspectable, unlike a latent code. Byte efficiency is a separate property and is evaluated independently through the codec comparison.

4.2. Byte-Level Capability Relative to ZFP and SZ3

Figure 11 places the adaptive reference beside calibration sweeps of ZFP and SZ3 applied to dense float32 fields after nearest-valid filling of invalid cells. The horizontal axis is the measured serialized byte ratio B, not the structural tile ratio. At the proposed calibration reference, the maximum adaptive payload ratio was 0.3362. The selected fidelity-constrained ZFP point had maximum $B = 0.1334$, RMSE 0.0525, normalized semivariogram discrepancy 0.0064, and physical-gradient magnitude loss 0.1232. The selected fidelity-constrained SZ3 point had maximum $B = 0.0318$, RMSE 0.0834, normalized semivariogram discrepancy 0.0075, and physical-gradient magnitude loss 0.1387. Both met the study-specific calibration fidelity limits with smaller payloads under this fill protocol.

This ordering is consistent with the codec designs. ZFP uses compact transformed block representations [27], whereas SZ3 composes prediction-based error-bounded modules [28,29]. The adaptive record additionally stores explicit topology and membership and therefore has a different payload objective. The multiple fidelity metrics follow the quantity-of-interest principle that compression quality should be assessed against downstream scientific structure as well as a single pointwise norm [31]. The result also accords with information-content studies that treat size, precision, and speed as distinct operational dimensions [32].

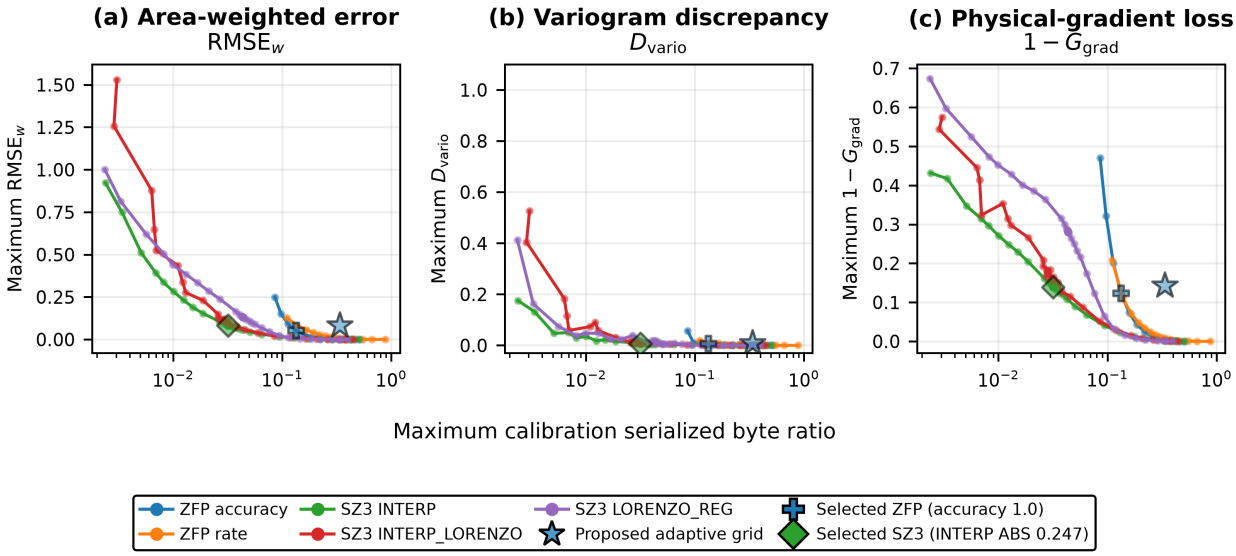


Figure 11. Calibration byte–fidelity operating regions for direct scientific codecs. Each panel plots maximum calibration serialized byte ratio against maximum (**a**) RMSE$_w$, (**b**) D_{vario}, or (**c**) $1 - G_{\text{grad}}$. Curves show tested direct ZFP and SZ3 candidates. The star is the selected adaptive-grid reference and is not a member of either codec curve; the black-edged filled-plus and diamond markers identify the selected fidelity-constrained ZFP and SZ3 operating points, respectively. All operating-point selections use calibration fields only. ABS denotes the absolute-error bound; INTERP, INTERP_LORENZO, and LORENZO_REG identify the SZ3 interpolation, interpolation/Lorenzo, and Lorenzo/regression configurations, respectively.

The complete canonical candidate sweep used to regenerate Figure 11 is provided as machine-readable Supplementary Data, together with the field-level measurements, calibration maxima, fixed candidate grids, byte-wrapper definition, and candidate definitions used to select the reported operating points.

4.3. Primary Held-Out Codec and Hybrid Comparison

Held-out results reinforce the distinction between numerical coding and explicit topology (Table 5; Figure 12). The proposed median payload ratio was $B = 0.3121$. Using medians of field-paired percentage changes, ZFP reduced payload by 62.7% and had lower values for all three losses. SZ3 reduced payload by 89.8%; its paired median changes were a 14.2% increase in RMSE, a 17.2% increase in normalized semivariogram discrepancy, and a 3.84% decrease in physical-gradient magnitude loss. SZ3 reduced semivariogram discrepancy in three of the eight fields, despite having a smaller separately aggregated median in Table 5. Relative summaries are therefore calculated before aggregation so that each method is compared with the proposed representation of the same field. Calibration fidelity matching transfers differently across codecs and metrics and is conditional on the stated fill protocol.

AG+ZFP and AG+SZ3 retain each field's adaptive partition and apply a codec only to its representative-value stream. Their median payload ratios were 0.1907 and 0.1725, with median field-paired reductions of 38.8% and 46.0%, respectively. Their reconstruction-loss summaries remained close to the proposed values, with all three paired median percentage changes below 1% in magnitude. These near-fidelity operating points show that explicit partitioning and mature value coding can be combined. Alternative hybrid operating points and their calibration trade-offs are reported in Appendix E.

This complementarity parallels prior work that treats resolution and precision as joint representation axes [25,26]. The present implementation is narrower: it uses a curvilinear masked domain, piecewise-constant tile representatives, and one field-specific topology rather than a general progressive multilinear representation. The comparison is therefore conceptual and functional.

Table 5. Primary held-out codec and hybrid comparison. Values are medians with 95% field-bootstrap intervals. Percentage changes are medians of field-paired changes relative to the proposed adaptive grid; negative values denote reductions.

Method (Selected Calibration Parameter)	B	RMSE_w	D_{vario}	$1 - G_{\mathrm{grad}}$	ΔB (%)
Proposed adaptive grid ($P_{\max} = 2$, $S_{\mathrm{th}} = 0.9737439$)	0.3121 [0.2968, 0.3481]	0.0708 [0.0653, 0.0760]	0.0051 [0.0029, 0.0059]	0.1304 [0.1210, 0.1389]	0
ZFP fidelity-constrained (fixed accuracy 1.0)	0.1284 [0.1139, 0.1347]	0.0508 [0.0501, 0.0514]	0.0034 [0.0022, 0.0044]	0.1018 [0.0971, 0.1105]	−62.7
SZ3 fidelity-constrained (INTERP ABS 0.247)	0.0313 [0.0257, 0.0383]	0.0815 [0.0797, 0.0820]	0.0048 [0.0034, 0.0086]	0.1221 [0.1150, 0.1286]	−89.8
AG+ZFP, tile-mean tolerance 0.005	0.1907 [0.1788, 0.2126]	0.0708 [0.0653, 0.0763]	0.0051 [0.0029, 0.0059]	0.1305 [0.1210, 0.1390]	−38.8
AG+SZ3, LORENZO_REG ABS 0.005	0.1725 [0.1609, 0.1862]	0.0709 [0.0653, 0.0764]	0.0051 [0.0032, 0.0059]	0.1310 [0.1214, 0.1384]	−46.0

Table 5 reports achieved held-out outcomes rather than exact equality constraints. Codec settings were chosen from calibration fields, so their held-out medians need not match the adaptive reference in any one metric; the intervals summarize variation across the eight sampled fields. Payload reductions should therefore be read together with the direction and magnitude of all three fidelity changes.

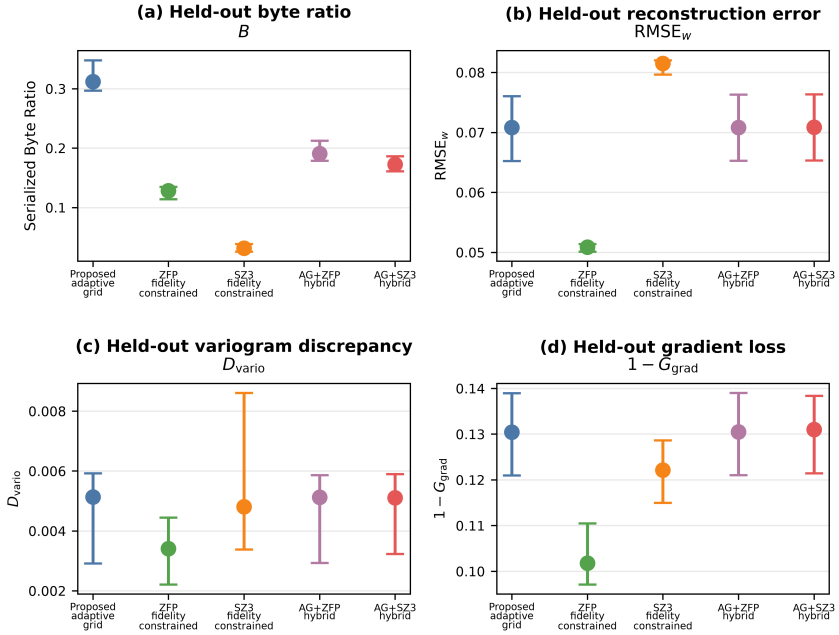


Figure 12. Primary held-out comparison of the proposed representation, direct codecs, and near-fidelity hybrids. Points are eight-field medians and bars are 95% field-bootstrap intervals for (**a**) serialized byte ratio, (**b**) area-weighted RMSE, (**c**) normalized semivariogram discrepancy, and (**d**) physical-gradient magnitude loss. AG+ZFP and AG+SZ3 retain the adaptive partition and compress only representative tile values.

Learned marine and geospatial compression provides another important comparison class [33,36]. Fixed-quality neural coding and learned spectral–spatial transforms for remote-sensing imagery illustrate recent progress toward controlled quality and cross-band representation [34,35]. A direct benchmark was not undertaken because training data, architectures, rate definitions, and target dimensionality require a separate controlled design. Future comparisons should measure distribution shift, training and inference cost, task-aware fidelity, coordinate support, and direct spatial-query capability.

4.4. Operational Mapping and Implementation Cost

An encoded adaptive record carries an anchored row–column origin, dyadic level or side, representative value, and a schema from which source-cell membership is reconstructed using the shared array shape and validity mask. A prospective client could select tile records intersecting a geographic region through the shared coordinates, reconstruct relevant cells, or retrieve the authoritative source for high-precision analysis. This study implements deterministic full-field reconstruction but does not benchmark a regional-query service. Common coordinates and mask are shared reference assets excluded from the primary payload ratios; a service that cannot amortize them must include their transfer and versioning costs.

Appendix D separates timing boundaries that were previously conflated. Median one-shot encode times over held-out fields were 696.25 ms for the proposed method, 2.63 ms for ZFP, 6.21 ms for SZ3, 694.41 ms for AG+ZFP, and 699.87 ms for AG+SZ3. Prepared-representation encode times were 13.29, 2.15, 5.76, 11.18, and 16.67 ms, respectively. Common-output decode times were 36.05 ms for the proposed representation, 1.35 ms for ZFP, 2.86 ms for SZ3, 36.51 ms for AG+ZFP, and 38.64 ms for AG+SZ3.

These values expose implementation cost but do not support an algorithm-independent speed claim. The compiled dense codecs and Python adaptive pipeline do not have matched native implementations, and metadata handling is intrinsic to the adaptive output. NetCDF I/O, candidate calibration, and metric computation are excluded. Supplementary Figure S7 further reports merge-core workload scaling; the theoretical traversal bound remains $O(n_r n_c P_{\max})$. A production comparison should harmonize language, compiler, threading, memory allocation, I/O, and representative selective-query workloads.

4.5. Limitations and Generalization

Six limitations determine the interpretation. First, the primary East Sea topology analysis uses twelve nominal monthly first-day fields identified by filename order; exact year, averaging interval, and forecast lead are not resolved by the retained files. Intervening daily, event-scale, synoptic, and within-month variation is not observed. The field bootstrap quantifies sensitivity to the eight sampled held-out fields rather than independent-population uncertainty.

Second, the same-grid salinity analysis and California Current VIIRS chlorophyll-a case study broaden the variable, region, grid, and observation-source scope but remain two finite examples. Direct transfer of the temperature profile differs markedly from target-specific calibration, and the external case does not establish universal parameter transfer. Boundary-limited semivariogram fits in the validation cases do not provide resolved correlation lengths or validate their caps. Third, the primary East Sea cap is nonbinding across the evaluated candidate set. The tighter-cap experiment verifies mechanical enforcement, but neither it nor the primary result establishes a fidelity benefit from the correlation-derived safeguard or validates $\alpha = 0.7$ as a physical constant.

Fourth, local range is a simple split statistic. The unconstrained-range quadtree gives similar aggregate losses, but it changes both maximum level and threshold and changes the containing tile for a substantial fraction of cells; it is not a component-isolating ablation. Exact tile differences are conditional on the array anchoring and do not by themselves establish a practical downstream benefit. The evidence supports the deterministic representation and calibration protocol rather than uniform fidelity superiority. Fifth, the selected threshold is conditional on the finite candidate set, the objective normalization, and the spatial-metric sampling design. The semivariogram uses 500 sampled cells and broad distance bins, so its magnitude should not be interpreted as a converged description of every local spatial scale. Its omnidirectional range also averages directional anisotropy; a direction-aware rule would require separately validated range estimates.

Sixth, the model-mask-boundary experiment is a single M01 mechanism test with a prescribed allocation prior, not an independently validated coastal model. Its distance field follows the source model's valid–invalid interface, which depends on mask construction and is not equivalent to a physical shoreline. At $\rho = 0.60$ and fixed tile count, the nearest-quartile loss reductions are paid for by higher farthest-quartile losses and by increases in global RMSE and gradient loss; only the global semivariogram discrepancy improves. The same-base-tolerance case improves all three global losses only by using 17.74% more tiles. These outcomes demonstrate controllable redistribution, but they neither identify an optimal strength nor establish benefit for other fields, masks, coastlines, variables, or downstream tasks. Such claims require a prespecified application objective and held-out validation of both the spatial priority and its budget–fidelity consequences.

The multivariate, vector, and three-dimensional formulas in the Supplementary Materials are candidate extensions that remain to be validated. Neural compression and content-adaptive remote-sensing encoders [24,33–36] likewise remain outside the controlled benchmark. Recent global ocean-forecasting systems highlight the further need to test representations on multivariate, dynamically evolving fields and with process-oriented diagnostics [9,10].

5. Conclusions

This study developed and evaluated a deterministic, constraint-based adaptive representation for two-dimensional scalar marine fields. The method combines globally anchored dyadic merging, a field-normalized local-range tolerance, means weighted by center-derived polygon areas, a WGS84 realized-tile diameter safeguard, and lossless tile-to-cell membership. One parameter profile was selected within an evaluated candidate set using four calibration fields, four-objective Pareto filtering, and normalized minimax regret. The profile was then fixed while a new partition was generated from each of eight within-sequence held-out nominal monthly first-day fields.

Within the East Sea temperature experiment, the fixed profile retained a median 20.49% of valid cells as explicit tiles while maintaining median area-weighted RMSE 0.0708, normalized semivariogram discrepancy 0.0051, and physical-gradient magnitude loss 0.1304. It had lower losses than the strict-mask nearest-count uniform baseline, but comparison with fixed 2×2 and 4×4 references showed metric-dependent compactness–fidelity tradeoffs. The unconstrained-range quadtree produced similar aggregate losses with different tile boundaries; because both its maximum level and threshold differ, it is a comparator rather than a component ablation. The correlation-derived cap was nonbinding in the primary candidate set. Tighter-cap diagnostic tests constrained merges and altered tile allocation, confirming implementation without demonstrating that the primary cap caused the observed performance.

The deterministic framework was also demonstrated for same-grid salinity and external NOAA VIIRS chlorophyll-a case studies. Direct temperature-profile transfer without recalibration and target-specific calibration produced different compactness–fidelity balances. The external chlorophyll semivariogram range was unresolved at the fit bound and therefore does not validate the scale cap.

Direct ZFP and SZ3 achieved greater serialized payload reduction under the stated float32 benchmark and dense-fill protocol. The proposed method contributes an explicit, scale-admissible spatial partition with exact tile-to-cell membership. Hybrid experiments reduced the adaptive payload while retaining closely similar reconstruction-loss metrics, demonstrating that explicit topology and value coding can be combined.

The conclusions remain restricted to the tested monthly sequences and case studies. Only one nominal first-day East Sea field per month was analyzed, so daily and within-month topology remain untested. Broader claims require additional regions, variables, depths, resolutions, sensing products, active-cap cases, and high-frequency sequences, together with native implementation benchmarks. Under these boundaries, the proposed method offers a transparent derived layer intended for multiresolution visualization, selective transfer, and exploratory marine-data services while leaving the authoritative numerical or observational grid unchanged.

Supplementary Materials: The following supporting information can be downloaded at https://www.mdpi.com/article/10.3390/rs18193349/s1, Figures S1–S7, Tables S1–S11, losslessly compressed East Sea analysis-variable subsets, the regional VIIRS NetCDF input, machine-readable scientific comma-separated values (CSV) tables, source code, software-version and dependency records, and the executed analysis notebook.

Author Contributions: Conceptualization, J.P.; methodology, J.P.; software, J.P.; validation, J.P. and S.Y.K.; formal analysis, J.P.; investigation, J.P.; data curation, J.P.; writing—original draft preparation, J.P.; writing—review and editing, J.P. and S.Y.K.; visualization, J.P.; supervision, S.Y.K.; project administration, J.P. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Ministry of Oceans and Fisheries, Republic of Korea, through A Research for building and utilizing marine space digital twin models (grant no. RS-2022-KS221625).

Data Availability Statement: The ROMS numerical-prediction data used for the East Sea temperature and salinity analyses are public data provided by the Korea Hydrographic and Oceanographic Agency (KHOA) through the Korean Public Data Portal [38]. They may be used and reused under the public-work free-use/open-use license terms indicated for the dataset, subject to source attribution and other applicable conditions. The NOAA S-NPP VIIRS science-quality monthly chlorophyll-a product is publicly available from NOAA CoastWatch through the Environmental Research Division Data Access Program (ERDDAP) [39]. The accompanying Supplementary Data include twelve losslessly compressed East Sea NetCDF subsets containing `temp`, `salt`, `lon`, `lat`, and `mask`. These subsets preserve the retained arrays at their original float64 precision but are not copies of the complete source files. The regional VIIRS NetCDF input used in this study is also included. Scientific CSV tables, source code, software-version and dependency records, the executed analysis notebook, the ordered input inventory, calibration–held-out assignments, and derived numerical artifacts accompany these inputs.

Acknowledgments: The authors gratefully acknowledge support from the Ministry of Oceans and Fisheries, Republic of Korea, through A Research for building and utilizing marine space digital twin models (grant no. RS-2022-KS221625). The authors also thank the Korea Hydrographic and Oceanographic Agency (KHOA) for making the ROMS numerical-prediction data publicly available, and NOAA's Center for Satellite Applications and Research and CoastWatch program for providing the chlorophyll-a data.

Conflicts of Interest: Author Jinhong Park was employed by Sfractum Co., Ltd. The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ABS	Absolute-error bound
AEQD	Azimuthal equidistant projection
AG	Adaptive grid
CSV	Comma-separated values
ERDDAP	Environmental Research Division Data Access Program
FESOM2	Finite-volumE Sea ice–Ocean Model
I/O	Input/output
KHOA	Korea Hydrographic and Oceanographic Agency
LAEA	Lambert azimuthal equal-area projection
MPAS-Ocean	Ocean component of the Model for Prediction Across Scales
NetCDF	Network Common Data Form
NOAA	National Oceanic and Atmospheric Administration
RMSE	Root mean square error
ROMS	Regional Oceanic Modeling System
S-NPP	Suomi National Polar-orbiting Partnership
VIIRS	Visible Infrared Imaging Radiometer Suite
WGS84	World Geodetic System 1984
SZ3	Modular prediction-based error-bounded scientific-data compressor
ZFP	Compressed-array representation for floating-point data

Appendix A. Algorithm and Reproducibility Details

For each field and parameter pair, the deterministic partition procedure is:

1. validate the finite ocean mask and initialize all valid cells as unassigned;
2. compute R_m and τ_m from Equation (3);
3. for $p = P_{\max}, P_{\max} - 1, \ldots, 1$, traverse globally anchored $2^p \times 2^p$ blocks in row-major order;
4. accept a block only when it is complete, fully valid, unassigned, scale admissible, and satisfies Equation (4);
5. assign every remaining valid cell to a singleton tile;
6. compute Equation (5), reconstruct the dense valid-ocean field, and verify exact coverage, no overlap, finite values, and deterministic membership;
7. on calibration fields only, evaluate Equations (8), (11) and (12), then apply Equations (13)–(15).

The reproducibility archive includes the lossless East Sea analysis-variable subsets and regional VIIRS input described in the Methods, the ordered input-file inventory, calibration–held-out assignments, candidate and comparator outputs, codec and hybrid summaries, source and executed notebooks, and analysis code for the reported structural, fidelity, payload, transfer, and sensitivity results. The primary environment was Python 3.12.10 with NumPy 2.5.2, pandas 3.0.5, xarray 2026.7.0, SciPy 1.18.1, Matplotlib 3.11.1, pyproj 3.7.2, pyshp 3.1.6, netCDF4 1.7.4, zfpy 1.0.1, and pysz 1.0.3. Exact software-version and dependency records, including optional-codec requirements, accompany the archive.

Appendix B. Physical-Coordinate and Spatial-Correlation Metrics

Figure A1 gives the scale diagnostic for M01. The selected $P_{max} = 2$ has a representative dyadic side of 13.370 km, while its exact maximum realized WGS84 corner-pair diameter is computed block by block from Equation (7). The fitted M01 practical range was 701.28 km. The fixed cap is not 0.7 times the M01 range alone; it is 0.7 times the minimum of the four calibration ranges and equals 366.9376 km. The choice $\alpha = 0.7$ was fixed before held-out evaluation as a prespecified design margin, not fitted as a coefficient or interpreted as a confidence bound. The maximum realized diameter was 21.157 km on the calibration set and 21.293 km across all twelve temperature fields, so the cap was nonbinding at the selected profile. Figure 7 reports the corresponding active-cap sensitivity.

The M01 variogram diagnostic used 110,160 retained sample pairs, a maximum lag of 1442.59 km, and 16 populated bins. The directional practical-range ratio was 5.256, indicating marked directional anisotropy, while the main selection retained the fixed omnidirectional estimate.

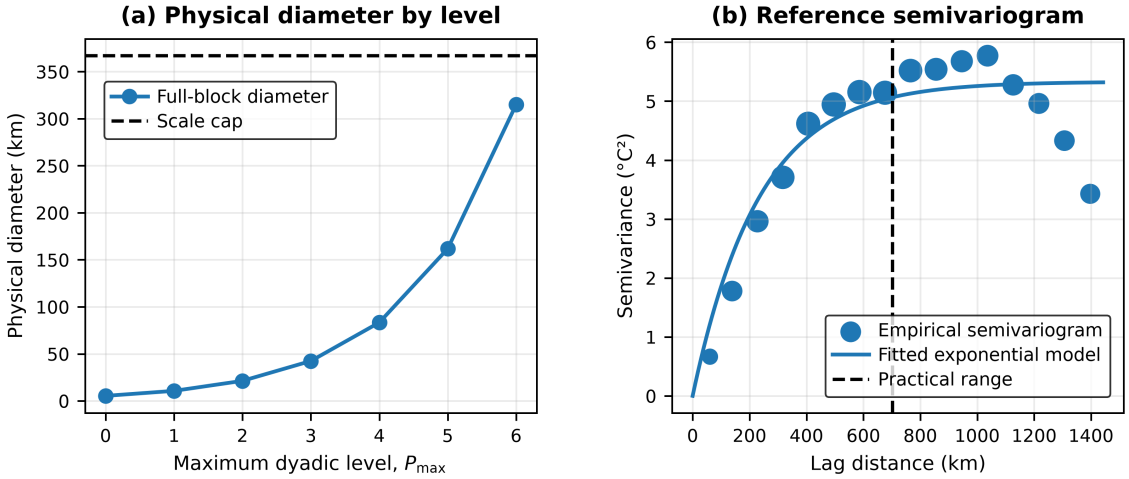


Figure A1. Physical metrics used by the adaptive-grid method. (**a**) Dyadic physical-scale diagnostics, including exact maximum WGS84 outer-corner diameter, together with the fixed 366.9376 km cap derived from the minimum calibration practical range. (**b**) M01 empirical semivariogram, pair-count-weighted exponential fit, and fitted practical range. Values provide an M01 reference diagnostic rather than a common range for all fields.

Appendix C. Alternative Codec Selection Rules

The primary comparison is only one operational question. A common-byte-budget sensitivity first retains candidates whose maximum calibration payload does not exceed the proposed maximum payload. Among these, direct codecs minimize the maximum absolute valid-ocean reconstruction error, whereas hybrids minimize the maximum absolute tile-representative coding error. These are method-specific selection quantities, not a common aggregate of the three plotted fidelity losses. The budget is an upper bound rather than an exact byte match. A second sensitivity builds a method-specific Pareto set in byte ratio and the same three fidelity losses, normalizes within that family, and applies minimax regret. Because normalization bounds differ, regret values cannot be used to rank families.

Table A1 defines every primary and sensitivity operating point, and Table A2 reports the corresponding calibration maxima on the common byte and fidelity metrics.

The three selection rules answer complementary operational questions. The fidelity-constrained rule asks how small a dense-field codec payload can be while all calibration losses remain within the proposed-profile maxima. The common-budget rule instead asks how much numerical fidelity can be obtained without exceeding the proposed maximum

serialized payload. The family-specific minimax rule balances payload and the three losses after normalization within each candidate family. It is therefore a within-family selection device rather than a common score for ranking unlike representations.

Table A1. Canonical compressor and hybrid operating-point definitions.

Scenario	Method	Calibration-Selected Parameter
Proposed reference	Proposed adaptive grid	$P_{\max} = 2$; $S_{\text{th}} = 0.9737439$; lossless tile records
Primary fidelity-constrained	SZ3 fidelity-constrained	INTERP; ABS = 0.247
Primary fidelity-constrained	ZFP fidelity-constrained	fixed accuracy; requested and effective tolerance = 1
Primary fixed near-fidelity hybrid	AG+SZ3 hybrid	fixed partition + SZ3 `LORENZO_REG` tile means; ABS = 0.005
Primary fixed near-fidelity hybrid	AG+ZFP hybrid	fixed partition + ZFP-coded tile means; accuracy = 0.005
Common-byte-budget	AG+SZ3 budget-selected	fixed partition + lossless SZ3 tile means
Common-byte-budget	AG+ZFP budget-selected	fixed partition + lossless ZFP tile means
Common-byte-budget	SZ3 budget-selected	`LORENZO_REG`; ABS = 5.5×10^{-4}
Common-byte-budget	ZFP budget-selected	fixed accuracy; requested = 0.04, effective = 0.0313
Method-specific four-objective minimax	AG+SZ3 minimax	fixed partition + SZ3 `LORENZO_REG` tile means; ABS = 0.08
Method-specific four-objective minimax	AG+ZFP minimax	fixed partition + ZFP-coded tile means; accuracy = 0.1
Method-specific four-objective minimax	SZ3 minimax	INTERP; ABS = 0.08
Method-specific four-objective minimax	ZFP minimax	fixed accuracy; requested = 0.32, effective = 0.25

Table A2. Calibration maxima for the reported adaptive-grid, direct-codec, and hybrid operating points.

Method	B	RMSE_w	D_{vario}	$1 - G_{\text{grad}}$
AG+SZ3 budget-selected	0.2941	0.0835	0.0078	0.1424
AG+SZ3 hybrid	0.1814	0.0836	0.0078	0.1428
AG+SZ3 minimax	0.1324	0.0955	0.0082	0.1887
AG+ZFP budget-selected	0.2903	0.0835	0.0078	0.1424
AG+ZFP hybrid	0.2080	0.0835	0.0078	0.1425
AG+ZFP minimax	0.1801	0.0842	0.0075	0.1452
Proposed adaptive grid	0.3362	0.0835	0.0078	0.1424
SZ3 budget-selected	0.3352	0.0003	2.417×10^{-5}	0.0009
SZ3 fidelity-constrained	0.0318	0.0834	0.0075	0.1387
SZ3 minimax	0.0645	0.0312	0.0035	0.0679
ZFP budget-selected	0.2965	0.0025	0.0002	0.0069
ZFP fidelity-constrained	0.1334	0.0525	0.0064	0.1232
ZFP minimax	0.1869	0.0163	0.0022	0.0424

The achieved coordinates in Table A2 are reported with the rule labels because neither constraint implies exact matching. A budget-selected candidate may leave part of the permitted payload unused, and two fidelity-constrained candidates may distribute their residual loss differently among the three metrics. Moreover, the direct codecs return dense numerical reconstructions, whereas the adaptive and hybrid formats also retain explicit tile geometry and membership. The held-out comparisons below therefore report every achieved coordinate and keep output capability distinct from numerical fidelity.

Figure A2 shows that direct-codec settings selected under the calibration payload ceiling retain substantially lower held-out reconstruction losses. For example, the held-out median RMSE was 0.0024 for the budget-selected ZFP point and 0.0003 for the budget-selected SZ3 point. The calibration ceiling does not constrain held-out bytes: the SZ3 setting had a median paired payload increase of 1.1% relative to the proposed representation on held-out fields. The output also lacks the explicit adaptive topology.

Table A3 reports held-out medians for the same operating points, whereas Table A4 adds field-bootstrap intervals and makes the uncertainty across the sampled held-out fields explicit.

Figure A3 reports family-specific minimax choices. The selected ZFP and SZ3 points have held-out median byte ratios 0.1818 and 0.0630, respectively, with RMSE 0.0159 and

0.0304. Hybrid minimax points save more payload than the primary fixed near-fidelity hybrids but accept larger reconstruction losses, illustrating why the selection question must accompany any reported operating point.

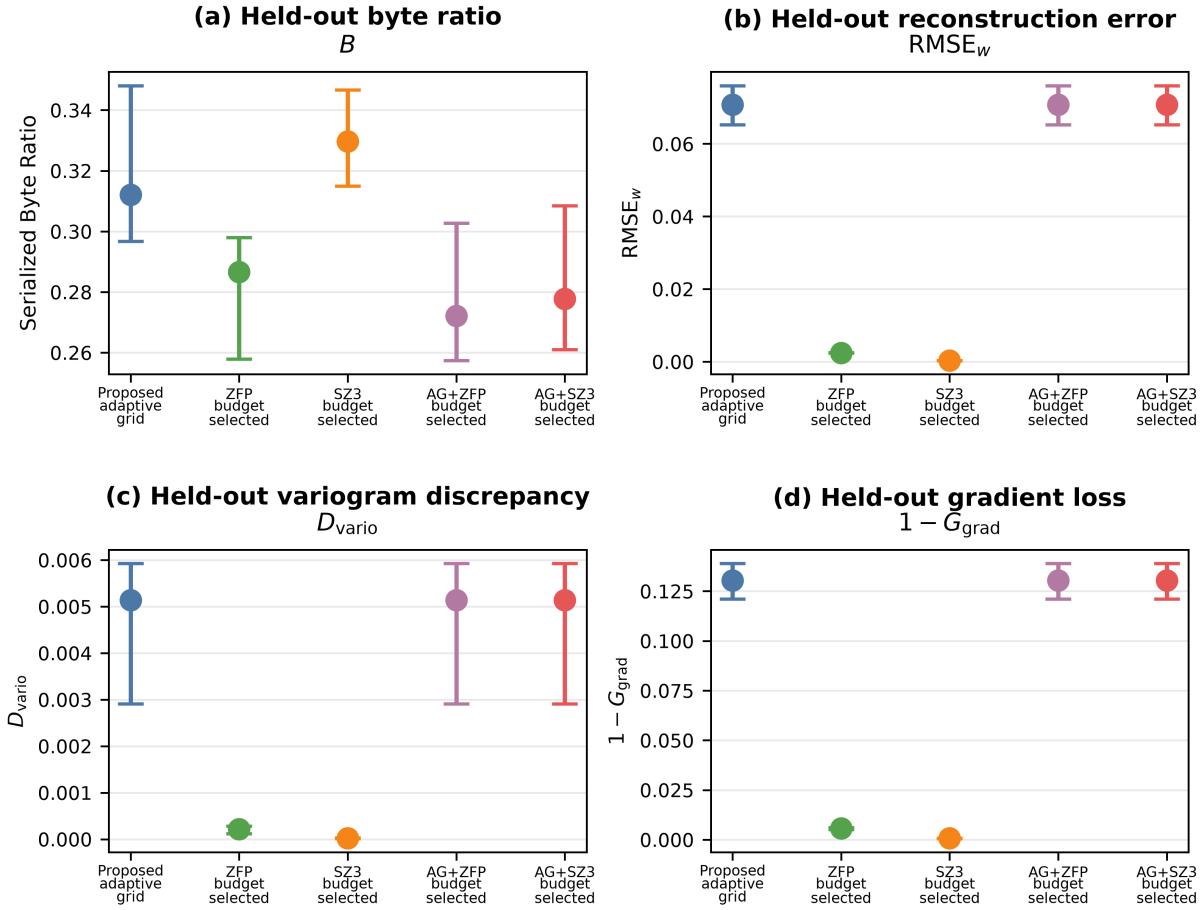


Figure A2. Held-out sensitivity under a common calibration maximum byte-budget rule. (**a**) Serialized byte ratio and (**b**–**d**) the three fidelity losses. A method may use less than the available budget; held-out byte ratios are not forced to be identical.

Table A3. Held-out medians for all reported adaptive-grid, direct-codec, and hybrid operating points.

Method	B	RMSE$_w$	D_{vario}	$1 - G_{\text{grad}}$
AG+SZ3 budget-selected	0.2777	0.0708	0.0051	0.1304
AG+SZ3 hybrid	0.1725	0.0709	0.0051	0.1310
AG+SZ3 minimax	0.1238	0.0855	0.0054	0.1685
AG+ZFP budget-selected	0.2721	0.0708	0.0051	0.1304
AG+ZFP hybrid	0.1907	0.0708	0.0051	0.1305
AG+ZFP minimax	0.1653	0.0715	0.0047	0.1333
Proposed adaptive grid	0.3121	0.0708	0.0051	0.1304
SZ3 budget-selected	0.3297	0.0003	2.294×10^{-5}	0.0007
SZ3 fidelity-constrained	0.0313	0.0815	0.0048	0.1221
SZ3 minimax	0.0630	0.0304	0.0029	0.0563
ZFP budget-selected	0.2867	0.0024	0.0002	0.0056
ZFP fidelity-constrained	0.1284	0.0508	0.0034	0.1018
ZFP minimax	0.1818	0.0159	0.0012	0.0342

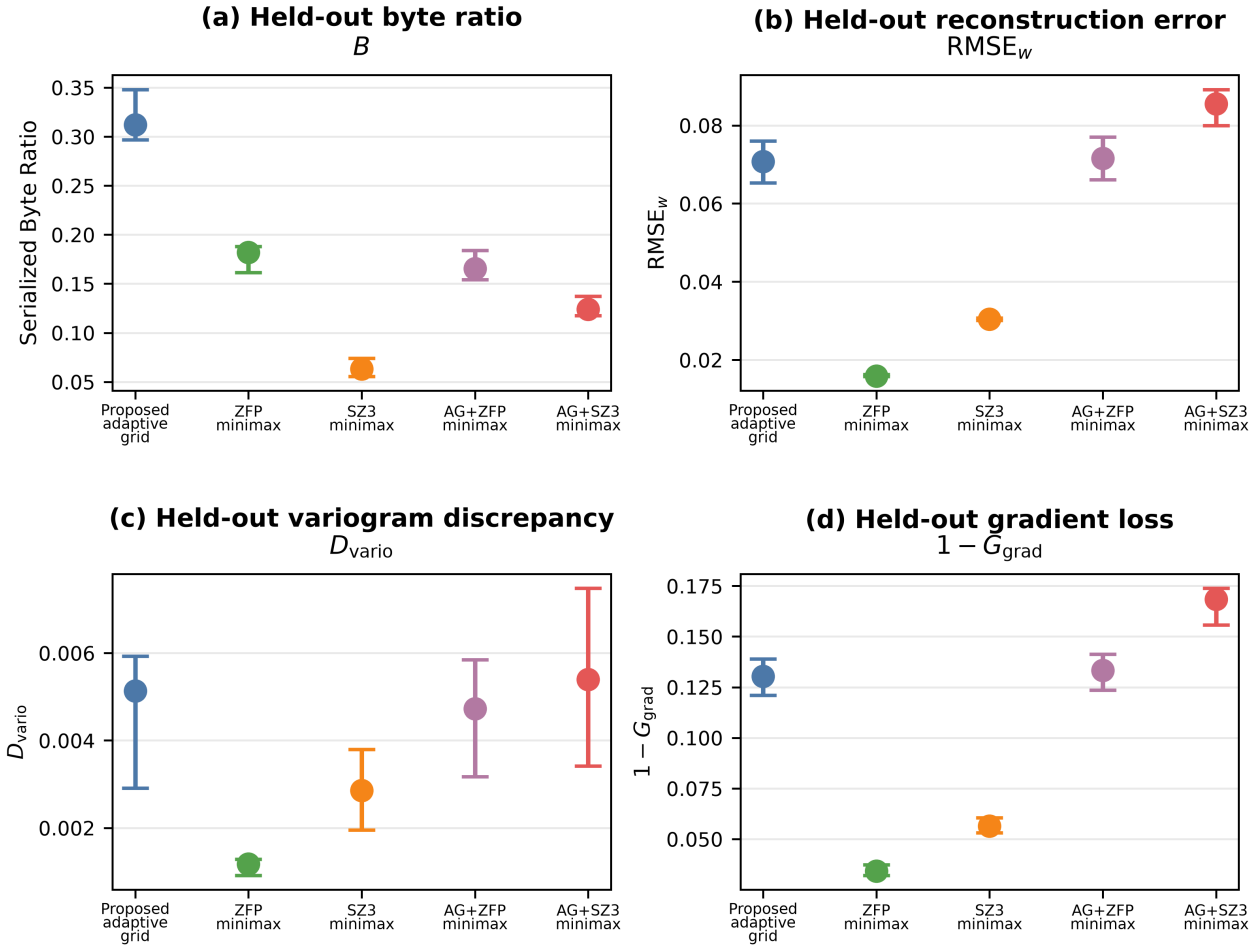


Figure A3. Held-out sensitivity for method-specific four-objective minimax selection. Each family is normalized using its own calibration Pareto extrema. The panel therefore compares the resulting operating points, not the family-specific normalized regret values themselves.

Table A4. Held-out medians and 95% field-bootstrap intervals for all reported operating points (5000 resamples). Each metric entry is the median [2.5th, 97.5th percentile].

Method	B	RMSE$_w$	D_{vario}	$1 - G_{\text{grad}}$
AG+SZ3 budget-selected	0.2777 [0.2611, 0.3085]	0.0708 [0.0653, 0.0760]	0.0051 [0.0029, 0.0059]	0.1304 [0.1210, 0.1389]
AG+SZ3 hybrid	0.1725 [0.1609, 0.1862]	0.0709 [0.0653, 0.0764]	0.0051 [0.0032, 0.0059]	0.1310 [0.1214, 0.1384]
AG+SZ3 minimax	0.1238 [0.1173, 0.1373]	0.0855 [0.0800, 0.0892]	0.0054 [0.0034, 0.0075]	0.1685 [0.1557, 0.1739]
AG+ZFP budget-selected	0.2721 [0.2574, 0.3027]	0.0708 [0.0653, 0.0760]	0.0051 [0.0029, 0.0059]	0.1304 [0.1210, 0.1389]
AG+ZFP hybrid	0.1907 [0.1788, 0.2126]	0.0708 [0.0653, 0.0763]	0.0051 [0.0029, 0.0059]	0.1305 [0.1210, 0.1390]
AG+ZFP minimax	0.1653 [0.1543, 0.1839]	0.0715 [0.0661, 0.0770]	0.0047 [0.0032, 0.0059]	0.1333 [0.1235, 0.1413]
Proposed adaptive grid	0.3121 [0.2968, 0.3481]	0.0708 [0.0653, 0.0760]	0.0051 [0.0029, 0.0059]	0.1304 [0.1210, 0.1389]
SZ3 budget-selected	0.3297 [0.3150, 0.3467]	0.0003 [0.0003, 0.0003]	2.294×10^{-5} [1.628×10^{-5}, 2.743×10^{-5}]	0.0007 [0.0006, 0.0008]
SZ3 fidelity-constrained	0.0313 [0.0257, 0.0383]	0.0815 [0.0797, 0.0820]	0.0048 [0.0034, 0.0086]	0.1221 [0.1150, 0.1286]
SZ3 minimax	0.0630 [0.0555, 0.0741]	0.0304 [0.0301, 0.0307]	0.0029 [0.0020, 0.0038]	0.0563 [0.0531, 0.0605]
ZFP budget-selected	0.2867 [0.2579, 0.2979]	0.0024 [0.0024, 0.0025]	0.0002 [0.0001, 0.0003]	0.0056 [0.0051, 0.0060]
ZFP fidelity-constrained	0.1284 [0.1139, 0.1347]	0.0508 [0.0501, 0.0514]	0.0034 [0.0022, 0.0044]	0.1018 [0.0971, 0.1105]
ZFP minimax	0.1818 [0.1615, 0.1881]	0.0159 [0.0158, 0.0163]	0.0012 [0.0009, 0.0013]	0.0342 [0.0321, 0.0373]

Appendix D. Harmonized Runtime Protocol

Figure A4 reports the three timing boundaries defined in Section 2. The field-level medians were bootstrapped over the same eight held-out fields. One-shot adaptive timings include construction of the partition; prepared-representation timings remove that construction and expose serialization of an already prepared representation. Common-output decoding requires every method to return a dense float32 ocean reconstruction.

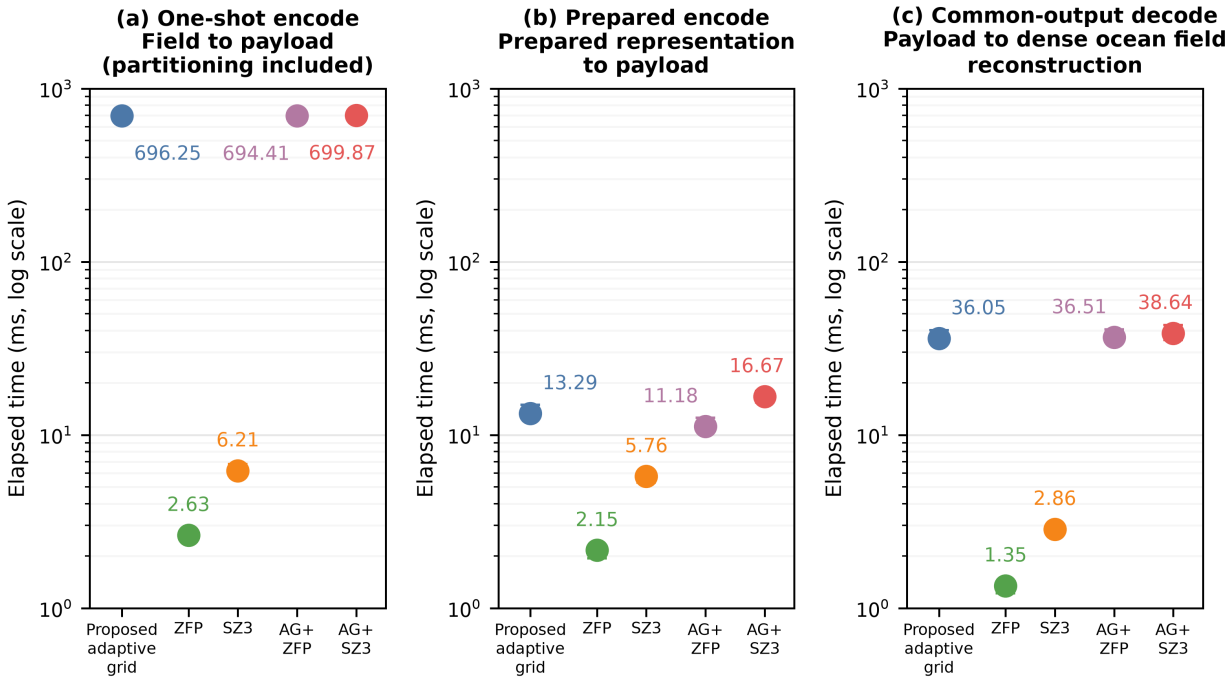


Figure A4. Harmonized implementation timing diagnostic. (**a**) One-shot field-to-payload encoding. (**b**) Prepared-representation-to-payload encoding. (**c**) Payload-to-dense-float32 common-output decoding. Points are held-out medians and bars are field-bootstrap intervals. NetCDF I/O, calibration, and fidelity metrics are excluded; compiled codec paths and Python adaptive paths prevent a language-neutral speed ranking.

Appendix E. Hybrid Operating-Space Analysis

Tables A5 and A6 report field-level metrics and held-out changes for the primary operating points. All reported byte ratios use the common denominator and measured wrapper rules defined in the Methods. The adaptive and hybrid records preserve the same tile geometry and membership scope but use their respective serialization layouts, whereas the direct codecs contain dense coded values and no adaptive-topology record.

The field-level rows separate effects of the fixed partition from effects of representative-value coding. Within each hybrid family, tile geometry, membership, and tile count are identical to the proposed representation; changes in payload arise from coding the representative-value stream, and changes in fidelity appear after those values are decoded and expanded to the dense ocean grid. Direct ZFP and SZ3 operate on the dense field and therefore do not share this fixed-topology interpretation.

Table A6 expresses medians of field-paired percentage changes relative to the proposed representation, rather than percentage changes between separately aggregated medians. Because all four reported quantities are minimized, a negative percentage denotes a reduction in payload or loss and a positive percentage denotes an increase. The signs are deliberately not collapsed into a single score: a smaller payload can accompany either smaller or larger fidelity losses, and the appropriate trade-off depends on the selection rule and the required output capability.

Figure A5 shows candidate projections after applying ZFP or SZ3 to tile representatives under a fixed adaptive partition. The primary choices keep calibration fidelity ratios close to the unhybridized adaptive reference, while the wider candidate sets expose the payload–fidelity alternatives available when representative values alone are coded.

Table A5. Field-level metrics for the five primary methods. Cal. denotes calibration; the complete 156-row machine-readable CSV file includes alternative selection scenarios.

Field	Role	B	RMSE_w	D_{vario}	$1 - G_{\mathrm{grad}}$
Proposed adaptive grid					
M01	Cal.	0.3236	0.0835	0.0066	0.1291
M02	Held-out	0.3069	0.0758	0.0037	0.1319
M03	Held-out	0.3020	0.0717	0.0054	0.1289
M04	Cal.	0.2949	0.0686	0.0032	0.1424
M05	Held-out	0.2944	0.0653	0.0049	0.1437
M06	Held-out	0.3173	0.0699	0.0056	0.1320
M07	Cal.	0.3362	0.0708	0.0078	0.1274
M08	Held-out	0.3894	0.0637	0.0059	0.0942
M09	Held-out	0.2968	0.0774	0.0029	0.1389
M10	Cal.	0.3146	0.0625	0.0032	0.1369
M11	Held-out	0.3402	0.0668	0.0102	0.1250
M12	Held-out	0.3481	0.0763	0.0028	0.1210
ZFP fidelity-constrained					
M01	Cal.	0.1191	0.0525	0.0029	0.0921
M02	Held-out	0.1139	0.0507	0.0032	0.0991
M03	Held-out	0.1110	0.0489	0.0022	0.1017
M04	Cal.	0.1089	0.0492	0.0035	0.1148
M05	Held-out	0.1216	0.0501	0.0021	0.1217
M06	Held-out	0.1336	0.0511	0.0041	0.1068
M07	Cal.	0.1334	0.0510	0.0028	0.1030
M08	Held-out	0.1402	0.0514	0.0053	0.0891
M09	Held-out	0.1347	0.0510	0.0037	0.1019
M10	Cal.	0.1243	0.0489	0.0064	0.1232
M11	Held-out	0.1268	0.0506	0.0044	0.1105
M12	Held-out	0.1299	0.0531	0.0026	0.0971
SZ3 fidelity-constrained					
M01	Cal.	0.0311	0.0834	0.0041	0.1112
M02	Held-out	0.0277	0.0819	0.0049	0.1195
M03	Held-out	0.0245	0.0797	0.0086	0.1247
M04	Cal.	0.0242	0.0793	0.0075	0.1292
M05	Held-out	0.0308	0.0793	0.0058	0.1378
M06	Held-out	0.0383	0.0820	0.0033	0.1251
M07	Cal.	0.0318	0.0809	0.0023	0.1182
M08	Held-out	0.0388	0.0818	0.0047	0.1015
M09	Held-out	0.0319	0.0812	0.0034	0.1150
M10	Cal.	0.0207	0.0795	0.0046	0.1387
M11	Held-out	0.0257	0.0805	0.0092	0.1286
M12	Held-out	0.0380	0.0842	0.0041	0.1167
AG+ZFP hybrid					
M01	Cal.	0.1940	0.0835	0.0066	0.1291
M02	Held-out	0.1825	0.0758	0.0038	0.1320
M03	Held-out	0.1788	0.0717	0.0054	0.1290
M04	Cal.	0.1753	0.0686	0.0031	0.1425
M05	Held-out	0.1781	0.0653	0.0049	0.1437
M06	Held-out	0.1946	0.0699	0.0056	0.1321
M07	Cal.	0.2080	0.0708	0.0078	0.1275
M08	Held-out	0.2423	0.0637	0.0059	0.0944
M09	Held-out	0.1869	0.0774	0.0029	0.1390
M10	Cal.	0.1967	0.0625	0.0032	0.1369
M11	Held-out	0.2106	0.0668	0.0103	0.1251
M12	Held-out	0.2126	0.0763	0.0028	0.1210

Table A5. *Cont.*

Field	Role	B	RMSE$_w$	D_{vario}	$1 - G_{\text{grad}}$
AG+SZ3 hybrid					
M01	Cal.	0.1778	0.0836	0.0066	0.1295
M02	Held-out	0.1691	0.0758	0.0036	0.1324
M03	Held-out	0.1609	0.0717	0.0057	0.1296
M04	Cal.	0.1588	0.0687	0.0031	0.1428
M05	Held-out	0.1606	0.0653	0.0047	0.1442
M06	Held-out	0.1743	0.0700	0.0055	0.1326
M07	Cal.	0.1814	0.0708	0.0078	0.1280
M08	Held-out	0.2083	0.0638	0.0059	0.0949
M09	Held-out	0.1707	0.0775	0.0031	0.1394
M10	Cal.	0.1684	0.0626	0.0034	0.1373
M11	Held-out	0.1790	0.0668	0.0101	0.1255
M12	Held-out	0.1862	0.0764	0.0028	0.1214

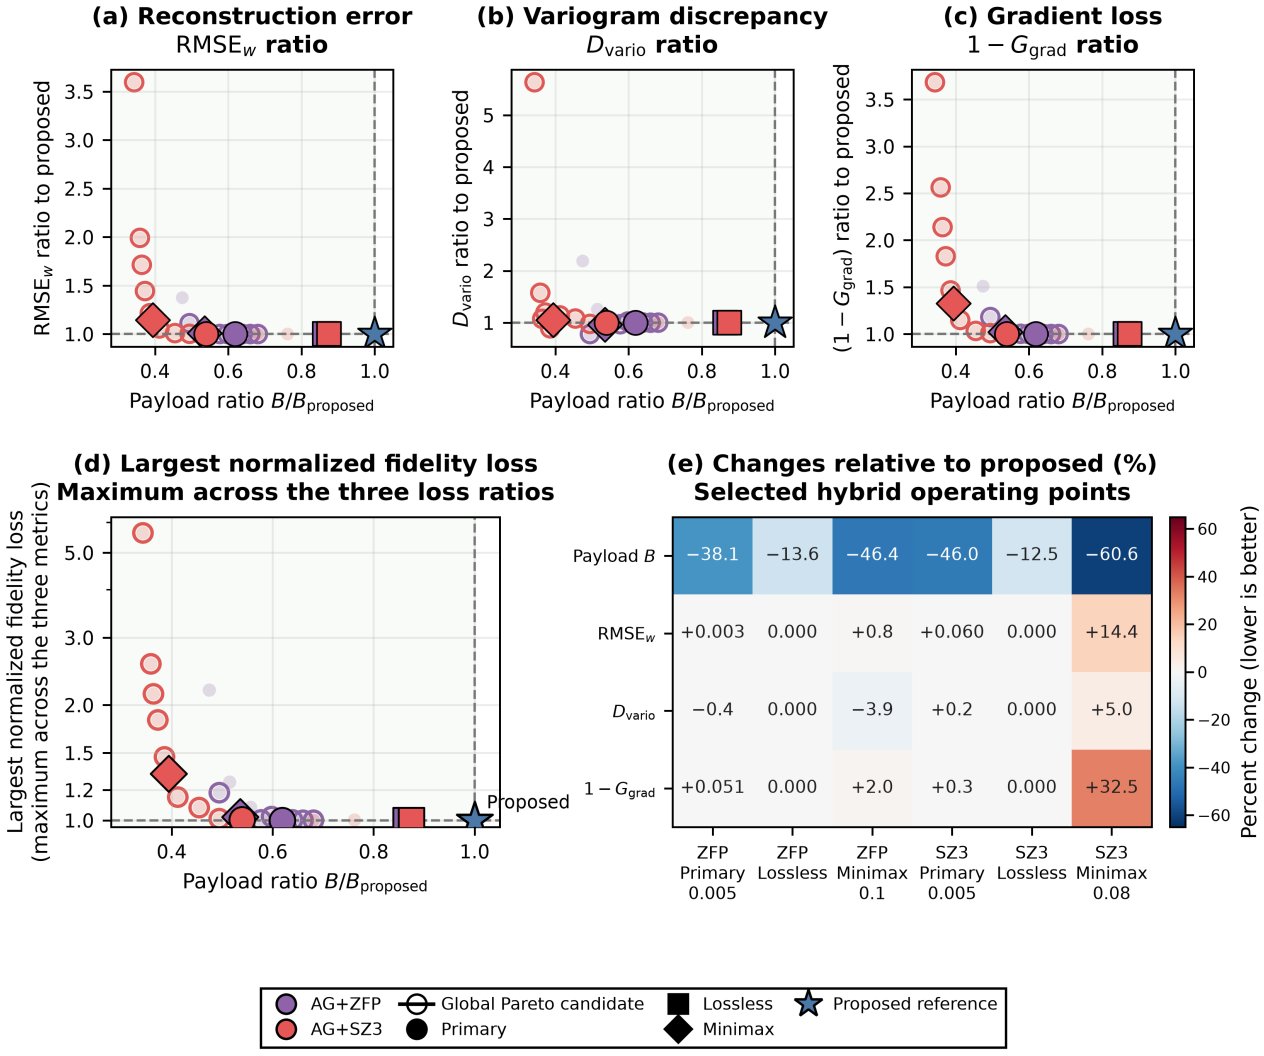


Figure A5. Hybrid operating-space analysis. (**a**–**c**) Normalized payload–loss projections for candidate AG+ZFP and AG+SZ3 tile-value settings. (**d**) Largest normalized fidelity loss. (**e**) Percentage changes at selected points. The partition, geometry, and membership remain fixed; only representative tile values are coded.

Table A6. Medians of held-out field-paired percentage changes relative to the proposed adaptive grid. Ratios are formed within each field before aggregation. All four numeric columns report percentages; negative values denote reductions.

Method/Scenario	ΔB	$\Delta \mathrm{RMSE}_w$	$\Delta D_{\mathrm{vario}}$	$\Delta(1-G_{\mathrm{grad}})$
AG+SZ3 budget-selected (Common-byte-budget)	−11.44	0	0	0
AG+SZ3 hybrid (Primary fixed near-fidelity hybrid)	−45.98	0.08	−0.47	0.40
AG+SZ3 minimax (Method-specific four-objective minimax)	−60.52	20.75	21.48	30.00
AG+ZFP budget-selected (Common-byte-budget)	−13.13	0	0	0
AG+ZFP hybrid (Primary fixed near-fidelity hybrid)	−38.81	0.00	0.12	0.06
AG+ZFP minimax (Method-specific four-objective minimax)	−47.04	1.05	−3.46	2.16
SZ3 budget-selected (Common-byte-budget)	1.11	−99.55	−99.57	−99.42
SZ3 fidelity-constrained (Primary fidelity-constrained)	−89.79	14.23	17.23	−3.84
SZ3 minimax (Method-specific four-objective minimax)	−79.89	−57.26	−44.90	−55.48
ZFP budget-selected (Common-byte-budget)	−15.98	−96.56	−96.26	−95.49
ZFP fidelity-constrained (Primary fidelity-constrained)	−62.70	−28.67	−21.40	−19.39
ZFP minimax (Method-specific four-objective minimax)	−47.08	−77.59	−76.16	−72.78

References

1. Woodcock, C.E.; Strahler, A.H. The factor of scale in remote sensing. *Remote Sens. Environ.* **1987**, *21*, 311–332. [CrossRef]
2. Lam, N.S.N.; Quattrochi, D.A. On the issues of scale, resolution, and fractal analysis in the mapping sciences. *Prof. Geogr.* **1992**, *44*, 88–98. [CrossRef]
3. Atkinson, P.M.; Tate, N.J. Spatial scale problems and geostatistical solutions: A review. *Prof. Geogr.* **2000**, *52*, 607–623. [CrossRef]
4. Grossmann, V.; Nakath, D.; Urlaub, M.; Oppelt, N.; Koch, R.; Köser, K. Digital twinning in the ocean—Challenges in multimodal sensing and multiscale fusion based on faithful visual models. *ISPRS Ann. Photogramm. Remote Sens. Spat. Inf. Sci.* **2022**, *4*, 345–352. [CrossRef]
5. Chen, G.; Yang, J.; Huang, B.; Ma, C.; Tian, F.; Ge, L.; Xia, L.; Li, J. Toward digital twin of the ocean: From digitalization to cloning. *Intell. Mar. Technol. Syst.* **2023**, *1*, 3. [CrossRef]
6. Miedtank, A.; Schneider, J.; Manss, C.; Zielinski, O. Marine digital twins for enhanced ocean understanding. *Remote Sens. Appl. Soc. Environ.* **2024**, *36*, 101268. [CrossRef]
7. Tzachor, A.; Hendel, O.; Richards, C.E. Digital twins: A stepping stone to achieve ocean sustainability? *npj Ocean Sustain.* **2023**, *2*, 16. Correction in *npj Ocean Sustain.* **2023**, *2*, 26. [CrossRef]
8. Banesh, D.; Petersen, M.R.; Ahrens, J.; Turton, T.L.; Samsel, F.; Schoonover, J.; Hamann, B. An image-based framework for ocean feature detection and analysis. *J. Geovis. Spat. Anal.* **2021**, *5*, 17. [CrossRef]
9. El Aouni, A.; Gaudel, Q.; Regnier, C.; Van Gennip, S.; Le Galloudec, O.; Drevillon, M.; Drillet, Y.; Lellouche, J.-M. GLONET: Mercator's end-to-end neural global ocean forecasting system. *J. Geophys. Res. Mach. Learn. Comput.* **2025**, *2*, e2025JH000686. [CrossRef]
10. Niu, Y.; Huang, Q.; Zhong, X.; Guo, A.; Chen, L.; Zhang, D.; Zhou, Z.-Q.; Jia, X.; Wu, L.; Zhang, R.; et al. A deep learning global ocean forecasting model with sub-daily and eddy-resolving resolution. *npj Clim. Atmos. Sci.* 2026, *epub ahead of printing*. [CrossRef]
11. Finkel, R.A.; Bentley, J.L. Quad trees: A data structure for retrieval on composite keys. *Acta Inform.* **1974**, *4*, 1–9. [CrossRef]
12. Samet, H. The quadtree and related hierarchical data structures. *ACM Comput. Surv.* **1984**, *16*, 187–260. [CrossRef]
13. Horowitz, S.L.; Pavlidis, T. Picture segmentation by a tree traversal algorithm. *J. ACM* **1976**, *23*, 368–388. [CrossRef]
14. Shusterman, E.; Feder, M. Image compression via improved quadtree decomposition algorithms. *IEEE Trans. Image Process.* **1994**, *3*, 207–215. [CrossRef] [PubMed]
15. Fu, G.; Zhao, H.; Li, C.; Shi, L. Segmentation for high-resolution optical remote sensing imagery using improved quadtree and region adjacency graph technique. *Remote Sens.* **2013**, *5*, 3259–3279. [CrossRef]
16. Toodesh, R.; Verhagen, S. Adaptive, variable resolution grids for bathymetric applications using a quadtree approach. *J. Appl. Geod.* **2018**, *12*, 311–322. [CrossRef]
17. Cong, Z.; Ma, T.; Li, Y.; Yuan, M.; Ling, Y.; Du, H.; Qi, C.; Li, Z.; Xu, S.; Zhang, Q. A storage-saving quadtree-based multibeam bathymetry map representation method. *J. Mar. Sci. Eng.* **2023**, *11*, 709. [CrossRef]
18. Li, M.; Stefanakis, E. Geospatial operations of discrete global grid systems—A comparison with traditional GIS. *J. Geovis. Spat. Anal.* **2020**, *4*, 26. [CrossRef]
19. Berger, M.J.; Oliger, J. Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **1984**, *53*, 484–512. [CrossRef]
20. Berger, M.J.; Colella, P. Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.* **1989**, *82*, 64–84. [CrossRef]
21. Popinet, S.; Rickard, G. A tree-based solver for adaptive ocean modelling. *Ocean Model.* **2007**, *16*, 224–249. [CrossRef]

22. Danilov, S.; Sidorenko, D.; Wang, Q.; Jung, T. The Finite-volumE Sea ice–Ocean Model (FESOM2). *Geosci. Model Dev.* **2017**, *10*, 765–789. [CrossRef]
23. Hoch, K.E.; Petersen, M.R.; Brus, S.R.; Engwirda, D.; Roberts, A.F.; Rosa, K.L.; Wolfram, P.J. MPAS-Ocean simulation quality for variable-resolution North American coastal meshes. *J. Adv. Model. Earth Syst.* **2020**, *12*, e2019MS001848. [CrossRef]
24. Hu, H.; Liu, C.; Liu, S.; Ying, S.; Wang, C.; Ding, Y. Full-process adaptive encoding and decoding framework for remote sensing images based on compression sensing. *Remote Sens.* **2024**, *16*, 1529. [CrossRef]
25. Hoang, D.; Summa, B.; Bhatia, H.; Lindstrom, P.; Klacansky, P.; Usher, W.; Bremer, P.-T.; Pascucci, V. Efficient and flexible hierarchical data layouts for a unified encoding of scalar field precision and resolution. *IEEE Trans. Vis. Comput. Graph.* **2021**, *27*, 603–613. [CrossRef] [PubMed]
26. Bhatia, H.; Hoang, D.; Morrical, N.; Pascucci, V.; Bremer, P.-T.; Lindstrom, P. AMM: Adaptive multilinear meshes. *IEEE Trans. Vis. Comput. Graph.* **2022**, *28*, 2350–2363. [CrossRef] [PubMed]
27. Lindstrom, P. Fixed-rate compressed floating-point arrays. *IEEE Trans. Vis. Comput. Graph.* **2014**, *20*, 2674–2683. [CrossRef] [PubMed]
28. Di, S.; Cappello, F. Fast error-bounded lossy HPC data compression with SZ. In Proceedings of the 2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS), Chicago, IL, USA, 23–27 May 2016; pp. 730–739. [CrossRef]
29. Liang, X.; Zhao, K.; Di, S.; Li, S.; Underwood, R.; Gok, A.M.; Tian, J.; Deng, J.; Calhoun, J.C.; Tao, D.; et al. SZ3: A modular framework for composing prediction-based error-bounded lossy compressors. *IEEE Trans. Big Data* **2023**, *9*, 485–498. [CrossRef]
30. Zhao, K.; Di, S.; Liang, X.; Li, S.; Tao, D.; Bessac, J.; Chen, Z.; Cappello, F. SDRBench: Scientific data reduction benchmark for lossy compressors. In Proceedings of the 2020 IEEE International Conference on Big Data, Atlanta, GA, USA, 10–13 December 2020; pp. 2716–2724. [CrossRef]
31. Jiao, P.; Di, S.; Guo, H.; Zhao, K.; Tian, J.; Tao, D.; Liang, X.; Cappello, F. Toward quantity-of-interest preserving lossy compression for scientific data. *Proc. VLDB Endow.* **2022**, *16*, 697–710. [CrossRef]
32. Klöwer, M.; Razinger, M.; Dominguez, J.J.; Düben, P.D.; Palmer, T.N. Compressing atmospheric data into its real information content. *Nat. Comput. Sci.* **2021**, *1*, 713–724. [CrossRef] [PubMed]
33. Gomes, C.; Wittmann, I.; Robert, D.; Jakubik, J.; Reichelt, T.; Maurogiovanni, S.; Vinge, R.; Hurst, J.; Scheurer, E.; Sedona, R.; et al. Lossy neural compression for geospatial analytics: A review. *IEEE Geosci. Remote Sens. Mag.* **2025**, *13*, 97–135. [CrossRef]
34. Mijares i Verdú, S.; Chabert, M.; Oberlin, T.; Serra-Sagristà, J. Fixed-quality compression of remote sensing images with neural networks. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2024**, *17*, 12169–12180. [CrossRef]
35. Mijares, S.; Bartrina-Rapesta, J.; Hernández-Cabronero, M.; Serra-Sagristà, J. Learned spectral and spatial transforms for multispectral remote sensing data compression. *IEEE Geosci. Remote Sens. Lett.* **2025**, *22*, 1–5. [CrossRef]
36. Sun, X.; Wang, P.; Zhou, Y.; Wu, K.; Huang, L.; Ma, X. Compression of marine environmental data using convolutional attention autoencoder. *J. Mar. Sci. Eng.* **2025**, *13*, 869. [CrossRef]
37. Shchepetkin, A.F.; McWilliams, J.C. The Regional Oceanic Modeling System (ROMS): A split-explicit, free-surface, topography-following-coordinate oceanic model. *Ocean Model.* **2005**, *9*, 347–404. [CrossRef]
38. Korea Hydrographic and Oceanographic Agency. ROMS Numerical Prediction Model Query Service. Korean Public Data Portal. 2025. Available online: https://www.data.go.kr/data/15142227/openapi.do (accessed on 4 September 2026).
39. NOAA CoastWatch. Chlorophyll, NOAA S-NPP VIIRS, Science Quality, Global 4 km, Level 3, 2012–Present, Monthly; Dataset Identifier `noaacwNPPVIIRSSQchlaMonthly`. Available online: https://coastwatch.noaa.gov/erddap/info/noaacwNPPVIIRSSQchlaMonthly/index.html (accessed on 6 September 2026).
40. Wang, M.; Son, S. VIIRS-derived chlorophyll-a using the ocean color index method. *Remote Sens. Environ.* **2016**, *182*, 141–149. [CrossRef]
41. Cressie, N.; Hawkins, D.M. Robust estimation of the variogram: I. *J. Int. Assoc. Math. Geol.* **1980**, *12*, 115–125. [CrossRef]
42. Matheron, G. Principles of geostatistics. *Econ. Geol.* **1963**, *58*, 1246–1266. [CrossRef]
43. Journel, A.G.; Huijbregts, C.J. *Mining Geostatistics*; Academic Press: London, UK, 1978.
44. Cressie, N. Fitting variogram models by weighted least squares. *Math. Geol.* **1985**, *17*, 563–586. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.